Contact Structures, σ-Confoliations, and Contaminations in 3-Manifolds

Ulrich Oertel [*] & Jacek Świątkowski [†]

Revised April, 2003

Rutgers University, Newark; Wrocław University, Wrocław

Abstract

We propose in this paper a method for studying contact structures in 3-manifolds by means of branched surfaces. We explain what it means for a contact structure to be carried by a branched surface embedded in a 3-manifold. To make the transition from contact structures to branched surfaces, we first define auxiliary objects called σ-confoliations and pure contaminations, both generalizing contact structures. We study various deformations of these objects and show that the σ-confoliations and pure contaminations obtained by suitably modifying a contact structure remember the contact structure up to isotopy.

After defining tightness for all pure contaminations in a natural way, generalizing the definition of tightness for contact structures, we obtain some conditions on (the embedding of) a branched surface in a 3-manifold sufficient to guarantee that any pure contamination carried by the branched surface is tight. We also find conditions sufficient to prove that a branched surface carries only overtwisted (non-tight) contact structures.

Our long-term goal in developing these methods is twofold: Not only do we want to study tight contact structures and pure contaminations, but we also wish to use them as tools for studying 3-manifold topology. These structures can exist in manifolds having either infinite or finite fundamental group, which suggests that they might be especially difficult to use effectively, but which also makes them especially attractive.

1 Introduction

We will usually assume that M is a closed orientable or oriented 3-manifold.

Definition 1.1. Suppose a field ξ of planes on an oriented 3-manifold M is locally defined by a form α. The field is integrable, and defines a *foliation* if $\alpha \wedge d\alpha = 0$. It defines a *positive contact structure* (*positive confoliation*) if $\alpha \wedge d\alpha > 0$ ($\alpha \wedge d\alpha \geq 0$). It defines a *negative contact structure* (*negative confoliation*) if $\alpha \wedge d\alpha < 0$ ($\alpha \wedge d\alpha \leq 0$). A positive (negative) confoliation is *contact at a point* if $\alpha \wedge d\alpha > 0$ ($\alpha \wedge d\alpha < 0$) at that point.

A (positive) contact structure is a field which is nowhere locally integrable due to the fact that the planes of the field locally look something like the tangent planes to the blades of a propeller (with the appropriate sense). A positive confoliation is a similar plane field, except that the limiting case of integrable fields is allowed, as is a mixture of integrable and positive non-integrable behavior. In this paper our working definition of a confoliation is often in terms of local charts called $\mathcal{B}$-charts and $\mathcal{C}$-charts, see Section 2.

The monograph [4] of Yakov Eliashberg and William Thurston made clear the close relationship between contact structures and foliations. If further evidence is needed, the recent papers [8, 9] confirm this close relationship. The work of Eliashberg and Thurston led us to consider the possibility of using branched surfaces to study contact structures. To make this approach possible, one must show that a contact structure can be modified such that it is "carried" by a branched surface, and one must further show that the modification can be done in such a way that the contact structure is not lost, or that it is recoverable from the modified object. As a first step in this direction we will show that if a contact structure is deformed slightly, to a special kind of confoliation called σ-confoliation, then the contact structure is indeed remembered.

Definition 1.2. A σ-*confoliation* is a confoliation ξ in a closed 3-manifold M for which there exists a smooth compact surface F in M such that:

(1) F is everywhere tangent to ξ,

(2) ξ is contact in the complement $M - F$,

(3) each connected component of F has nonempty boundary (equivalently, F admits a handle decomposition consisting of 0-handles and 1-handles only).

F is then called the *maximal integral surface* of the σ-confoliation ξ.

In the remainder of the paper, we will usually assume that all contact structures, confoliations, and contaminations (yet to be defined) are positive.

In studying modifications of contact structures into confoliations we distinguish the following classes of deformations.

Definition 1.3. A *deformation of contact structures* or *pure deformation* is a smooth deformation $\xi_t : t \in [a, b]$ of plane fields in M, such that for each $t \in [a, b]$ the plane field ξ_t is a contact structure. Similarly we define a *deformation of confoliations* as a smooth 1-parameter family of confoliations.

A deformation of confoliations (contact structures) ξ_t is an *isotopy of confoliations (contact structures)* if there is a smooth 1-parameter family ϕ_t of diffeomorphisms of M such that ϕ_0 is the identity and $(\phi_t)_*(\xi_0) = \xi_t$.

A σ-*deformation* is a smooth deformation $\xi_t : t \in [a, b]$ of plane fields in M, such that for each $t \in [a, b]$ the plane field ξ_t is a σ-confoliation, and moreover the maximal integral surfaces F_t of ξ_t behave well with respect to the parameter t in the following sense. There exists a 1-parameter family $h_t : t \in [a, b]$ of diffeomorphisms of M, a partition of the interval $[a, b]$ into a finite number of nontrivial disjoint consecutive subintervals $I_1, I_2, \ldots, I_k$, and a sequence $\Sigma_1, \Sigma_2, \ldots, \Sigma_k$ of compact smooth surfaces in M, such that:

(1) for each $t \in \mathrm{int}(I_j)$ we have $h_t(F_t) = \Sigma_j$;

(2) for each $1 \le j \le k - 1$ the surfaces Σ_j and Σ_{j+1} are distinct, and one of them is contained in the other;

(3) if $\Sigma_j \subset \Sigma_{j+1}$ (or conversely $\Sigma_{j+1} \subset \Sigma_j$), then the closure $cl(\Sigma_{j+1} - \Sigma_j)$ ($cl(\Sigma_j - \Sigma_{j+1})$ respectively) is a compact surface with piecewise smooth boundary;

(4) each component of the larger of the surfaces Σ_j and Σ_{j+1} is obtained from the smaller by attaching 1-handles and adding 0-handles.

In terms of these definitions, Gray's Theorem can be stated as follows: If ξ_0 and ξ_1 are connected by a deformation of contact structures, then they are isotopic.

In Sections 2, 5 and 6 we show how to construct various σ-deformations, both the ones that decrease and the ones that increase maximal integral surfaces. Our main result concerning σ-deformations is the following generalization of Gray's Theorem.

Theorem 1.4. σ-Deformation Theorem. *Any two contact structures on a closed 3-manifold M which are connected by a σ-deformation are isotopic.*

Theorem 1.4 implies that a contact structure is "remembered" or determined by a σ-confoliation connected to the contact structure by a σ-deformation. To see this, suppose that a σ-confoliation ξ is connected to contact structures ξ_0 and ξ_1 by σ-deformations. Combining the σ-deformations, we obtain a σ-deformation from ξ_0 to ξ_1, and therefore by the theorem ξ_0 and ξ_1 are isotopic.

The statement of Theorem 1.4 should be compared with Proposition 2.3.2 of [4], which implies that smooth deformations through arbitrary confoliations from one contact structure to another do not preserve the isomorphism class of the contact structures.

One of our most ambitious goals is to understand phenomenon of tightness of contact structures using σ-confoliations and branched surfaces. Towards the end of understanding tightness via σ-confoliations and branched surfaces, we give the following definition which extends the notion of tightness to all σ-confoliations.

Definition 1.5. A σ-confoliation ξ is *tight*, if there is no 2-dimensional disc D smoothly embedded in M, such that the boundary ∂D is tangent to ξ and disjoint from the maximal integral surface F of ξ, and D itself is transverse to ξ at ∂D.

Since we use Theorem 1.4 to identify a contact structure with a σ-confoliation if the two structures are related by a σ-deformation, the following theorem is needed to show that the above definition is consistent with the definition for contact structures.

Theorem 1.6. *Tightness of σ-confoliations is an invariant of σ-deformation.*

The statement of Theorem 1.6 should be compared with Proposition 3.6.2 of [4], which shows that smooth deformations through arbitrary confoliations do not preserve the property of tightness.

Now we introduce branched surfaces, and we explain what it means for a σ-confoliation to be carried by a branched surface. The following ideas, involving branched surfaces and contaminations will be explained more fully in Section 5.
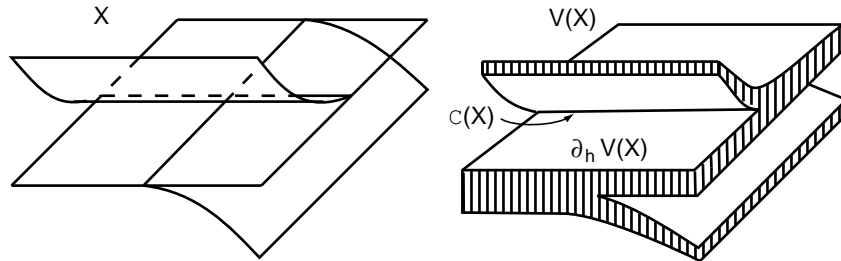

Figure 1: Local model for branched surface and fibered neighborhood.

A closed *branched surface B* with generic branch locus is a space with smooth structure locally modelled on the space X shown in Figure 1. (A neighborhood of each point of B is diffeomorphic to a neighborhood of a point in X.) The *branch locus K* of B is the union of points of B none of whose neighborhoods is a manifold. The branch locus is (homeomorphic to) a cell complex of dimension 1, but it can also be regarded as the image of the *branch curves $\mathcal{C}(B)$* immersed in B. The *sectors* of B are the completions (in a path metric coming from a Riemannian metric on B) of the components of $B - K$. The sense of branching at the boundary of a sector is described using the following convention. If sectors W, X, Y are adjacent along an arc γ of branch locus, and if $W \cup Y$ and $X \cup Y$ are smooth, we say that branching along the arc $\gamma \subset \partial Y$ is *inward* for Y and *outward* for X. We indicate the sense of branching using a transverse arrow pointing into Y. If B is embedded in a 3-manifold, then a *fibered neighborhood $V(B)$* of B in M is a closed regular neighborhood of B foliated by interval fibers, as shown in Figure 1, with the frontier $\partial V(B)$ cut by *cusp curves $\mathcal{C}(B)$* into smooth surfaces called the *components of horizontal boundary* and denoted $\partial_h V(B)$. The reader should not confuse $\partial_h V(B)$ with $\partial V(B)$, which is not smooth. We note that $\partial_h V(B)$ is not embedded in M or in $\partial V(B)$, and when we say that a subspace is contained in $\partial_h V(B)$ we mean that the embedding of the subspace in M factors through an embedding in

$\partial_h V(B)$. There is a projection map $\pi : V(B) \to V(B)/\sim$ which maps $V(B)$ to a quotient space in which fibers of $V(B)$ are collapsed to points. Notice that $V(B)/\sim$ can be identified with B, and that we identify the branch curves $\mathcal{C}(B)$ immersed in B with the cusp curves, also denoted $\mathcal{C}(B)$, which are embedded in $\partial V(B)$.

More detailed definitions of branched surfaces can be found in other papers, e.g. [7, 10, 11], or in the original source for branched manifolds, R. Williams, [13].

Definition 1.7. Suppose B is a branched surface embedded in a 3-manifold M. Suppose a smooth plane field ξ is defined on $V(B)$ such that:

(1) ξ is transverse to the interval fibers of $V(B)$;

(2) ξ is tangent to $\partial_h V(B)$;

(3) ξ is a confoliation in $V(B)$.
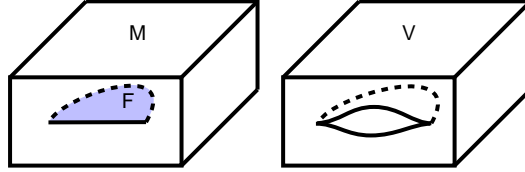
Then ξ is called a *contamination carried by B*.

Figure 2: Cutting a σ-confoliation.

Suppose now that ξ is a σ-confoliation in M. Then we can *cut M* on the maximal integral surface F to obtain V, see Figure 2, a *manifold with inward cusps*. This is a manifold with all the characteristics of a neighborhood $V(B)$ of a branched surface, except there is no foliation by interval fibers. The manifold has a smooth structure in a sense which will be made precise in Section 5. The boundary also has a smooth structure, but has inward cusps at a 1-submanifold $\mathcal{C}(V)$ of the boundary. We use $\partial_h V$, the *horizontal boundary*, to denote the components of ∂V cut on cusp curves. As before, we observe that $\partial_h V$ is not embedded in M, and that when we say that a subspace of M is contained in $\partial_h V(B)$ we mean that the embedding of the subspace in M factors through an embedding in $\partial_h V(B)$.

If the manifold with inward cusps V was obtained by cutting on F, then $\partial_h V$ consists of two copies of the integral surface F embedded in ∂V. In this case the manifold V can be embedded smoothly in M, such that the complement is topologically a product $\mathrm{int}(F) \times (0,1)$. The σ-confoliation ξ induces on V a smooth plane field, which we denote ξ'. This is the pure contamination in M induced by the σ-confoliation ξ. Regarding V as a subspace of M, the contamination ξ' is a partial plane field defined only on $V \subset M$. The operation converting ξ to ξ' is called *cutting ξ* on the maximal integral surface.

More generally:

5

Definition 1.8. A *(pure) contamination* in M is a smooth plane field ξ defined on a manifold V with inward cusps embedded smoothly in M, with the property that ξ is tangent to $\partial_h V$ and ξ is a (contact structure) confoliation in $\text{int}(V)$. The manifold V is called the *support* of the contamination ξ.

We note that just as a confoliation is a hybrid object obtained from foliations and contact structures, so a contamination is a hybrid object obtained from smooth laminations and contact structures.

We do not use arbitrary contaminations in this paper, only pure contaminations and certain other contaminations called σ-contaminations. In Section 5 we will give precise definitions of σ-contaminations, σ-deformations of σ-contaminations, and pure deformations of pure contaminations. Roughly speaking, a σ-contamination with support V is a contamination such that the integral surface consists of $\partial_h V$ together with a compact surface whose interior lies in the interior of V, and a σ-deformation is a deformation through σ-contaminations. Similarly, a pure deformation is a deformation through pure contaminations. It should be clear that σ-confoliations are σ-contaminations.

The operation of *splitting* a contact structure ξ converts the contact structure to a pure contamination ξ'' in two steps. First a σ-deformation converts ξ to a σ-confoliation ξ', then cutting on the maximal integral surface converts ξ' to a pure contamination ξ''.

We would like to represent an arbitrary contact structure ξ in M by a pure contamination obtained from ξ by means of splitting, with additional property that this contamination is carried by some branched surface $B \hookrightarrow M$. The following theorem shows this is possible.

Theorem 1.9. *Let ξ be a contact structure defined on a closed 3-manifold M and let Ω be a smooth 1-foliation transverse to ξ. There is a σ-deformation ξ_t, $0 \leq t \leq 1$, from $\xi = \xi_0$ to a σ-confoliation ξ_1 such that ξ_t is transverse to Ω at all times t and the maximal integral surface F for ξ_1 cuts all leaves of Ω into intervals. It follows that cutting ξ_1 on F yields a manifold V with inward cusps with a foliation by intervals induced by Ω, and thus V is isomorphic to the fibered neighborhood of a branched surface $B \hookrightarrow M$. Moreover, the resulting pure contamination supported by V is clearly carried by this branched surface and is a splitting of ξ_0.*

We would like to be able to recover a contact structure from a pure contamination obtained from the contact structure by a splitting. This can be done using a (roughly) opposite operation which we call pinching. Pinching can be applied to any pure contamination ξ in M supported by a manifold V satisfying the following condition: for any connected component U of $M - V$ the frontier ∂U consists of two homeomorphic components F of the horizontal boundary $\partial_h V$ (where $\partial F \neq \emptyset$) and U itself is homeomorphic to the product $\text{int}(F) \times (0, 1)$ (intuitively, U is homeomorphic to a thickened F with outward cusps at ∂F). If this condition is satisfied, we will say that contamination ξ

has *product complementary pieces.* Note that any pure contamination obtained by a splitting of a contact structure clearly has product complementary pieces.

Pinching a pure contamination ξ with product complementary pieces consists of the following steps (which will be described more precisely in Section 8). First, we identify the two components F in the frontier of any complementary component U via a diffeomorphism equal to the identity on ∂F (intuitively, we collapse U into a single integral surface F in M). We will call this operation a *glueing* and note that it depends not only on a glueing map, but also on the choice of smooth structure at the glueing locus (this will be discussed in detail in Section 8). Glueing gives us ξ', a sort of σ-confoliation, which is in general not smooth but only continuous (as a plane field in M) at the glueing locus, i.e. at the integral surfaces F obtained by collapsing the components U. We call such an object a *semi-smooth σ-confoliation.* We then modify ξ' to a contact structure, by simultaneously smoothing it and eliminating its maximal integral surface. This can be done by means of a *semi-smooth σ-deformation.* Semi-smooth σ-deformations will be defined in Section 8 and their existence will be proved there.

The operation inverse to a splitting is an example of a pinching, since we first reglue and then apply the reverse σ-deformation (which is of course a special case of a semi-smooth σ-deformation).

The following theorem shows that a contamination obtained from a contact structure by splitting remembers the contact structure.

Theorem 1.10. *(1) Let ξ be a pure contamination in a 3-manifold M with product complementary pieces. Then any two contact structures on M obtained from ξ by pinching are isotopic.*
(2) If ξ is a pure contamination obtained from a contact structure η on M by splitting (e.g. as in Theorem 1.9), then any contact structure obtained from ξ by pinching is isotopic to η.

The above theorem is a special case of a more general result, Lemma 8.6. In Section 8, we define splitting and pinching in the context of contaminations. Each splitting or pinching of a pure contamination yields another pure contamination. The splitting and pinching operations on pure contaminations generalize the splitting and pinching operations defined above for contact structures. We then prove that any two contact structures connected by any sequence of splittings and pinchings (of pure contaminations) are isotopic.

Definition 1.11. Let ξ be a contact structure on a 3-manifold M, and let B be a branched surface embedded in M. We say that B *carries* ξ if B carries a pure contamination obtained from ξ by a splitting.

It is not clear which branched surfaces B embedded in 3-manifolds carry contact structures. An obvious necessary condition is that B has product complementary pieces (which means that its fibered neighborhood $V(B)$ has product complementary pieces). In the separate paper [12], we give a sufficient condition for a branched surface to carry a positive pure contamination. Together

with the requirement of product complementary pieces, this gives a sufficient condition to carry a positive contact structure. On the other hand, a branched surface with product complementary pieces may carry two or more nonisotopic contact structures. Therefore in general one cannot recover a contact structure from a branched surface which carries it.

It turns out however, that one can impose conditions on a branched surface B so that any contact structure carried is overtwisted. We are able to distinguish such branched surfaces using "pseudo-transversal knots."

Definition 1.12. Let B be a branched surface embedded in M with product complementary pieces. A null-homologous knot κ in M is a *pseudo-transversal* for B if it satisfies the following conditions: (1) κ intersects B transversely away from the branch locus, (2) the intersection of κ with each complementary component U of $M - B$ is a collection of unknotted segments monotonically passing from one to the other component of the frontier ∂U (so that these segments form a braid in U).

It is possible to define a *self linking number* $l(\kappa)$ of a knot κ pseudo-transversal to B, in a way analogous to the self-linking (or the Thurston-Bennequin invariant) of a null-homologous knot transverse to a contact structure (see Section 10 for more details). If a branched surface B carries a contact structure ξ then κ determines (uniquely up to transverse isotopy) a knot κ' transverse to ξ, with the same self-linking. Transversal knots in tight contact structures are known to satisfy the so called Bennequin inequality

$$l(\kappa') \leq -\chi(S),$$

where $\chi(S)$ is the Euler characteristic of any Seifert surface S for κ', see [3]. By what was said above, the same inequality holds for pseudo-transversal knots κ, for any branched surface B that carries a tight contact structure. Thus we get the following.

Theorem 1.13. *Let B be a branched surface with product complementary pieces embedded in a 3-manifold M. Suppose that there is a knot in M which is pseudo-transversal to B and which violates the Bennequin inequality. Then any contact structure in M carried by B is overtwisted.*

It is clear that a branched surface B uniquely determines the homotopy type (as a plane field) of any contact structure it carries. In view of the fact that an overtwisted contact structure is uniquely determined (up to isotopy) by its homotopy type (see [2]), a branched surface B as in 1.13 carries at most one contact structure (up to isotopy) and this structure is overtwisted.

One of the most important goals of our research program is to prove results about 3-manifold topology using contact structures and contaminations. It is well-known that among contact structures on a 3-manifold, only tight ones can contain information about the topology of the 3-manifold. It is therefore necessary to find a suitable definition of tightness for pure contaminations. The

definition proposed by Eliashberg and Thurston for tightness of confoliations gives a clear indication for defining tightness of pure contaminations:

Definition 1.14. Suppose ξ supported by $V \hookrightarrow M$ is a pure contamination in M. Then a disc $D \hookrightarrow M$ is an *overtwisting disc* for ξ if the boundary ∂D is tangent to ξ, D itself is transverse to ξ at ∂D, and in case $\partial D \subset \partial_h V$, ∂D does not bound a disc in $\partial_h V(B)$. The pure contamination ξ is *tight* if there are no overtwisting discs for ξ.

We will show that in case the pure contamination ξ represents a contact structure, this definition implies tightness according to previous definitions. There are fairly obvious candidate definitions for tightness of arbitrary contaminations, but in this paper we restrict our attention to pure contaminations.

If we want to use tight contaminations to prove results about 3-manifolds, it will be necessary to understand modifications which preserve the property of tightness. Our knowledge in this domain is currently limited, but what we know is described in Section 11.

In Sections 13 and 14 we prove the "internal tightness" of some special classes of pure contaminations carried by branched surfaces. If ξ is a pure contamination carried by B, then ξ is *internally (universally) tight* if there are no overtwisting discs in $\text{int}(V(B))$ $(\text{int}(\tilde{V}(B)))$, where $\tilde{V}(B)$ is the universal cover of $V(B)$.

Theorem 1.15. *Suppose a branched surface $B \hookrightarrow M$ carries a positive σ-contamination ξ. Suppose that $\partial_h V(B)$ has no sphere components and suppose that the branch locus of B consists of a simple system of essential closed curves, i.e. there are no double points of the branch locus, and the smooth closed curves of the branch locus are π_1-injective. Then ξ is internally universally tight.*

A somewhat stronger version of this theorem (Theorem 13.3) is proved in Section 13, which has the following corollary. The definition of a "immersed twisted disc of contact" is given in Section 13, but the following corollary applies, in particular, when there are no double points of the branch locus.

Corollary 1.16. *Suppose that $B \hookrightarrow M$ is a branched surface satisfying the following conditions:*

(1) the branch locus of B consists of a system of essential closed curves;

(2) there are no immersed twisted discs of contact in B;

(3) the boundary $\partial V(B)$ is incompressible in M;

(4) $\partial_h V(B)$ does not contain sphere components or disc components.

Then any pure contamination carried by B is tight.

The above corollary can be used to construct examples of tight contaminations, see Section 14. There are conditions one can impose on a branched surface B sufficient to ensure than any contamination carried by B represents a tight contact structure.

9

It remains now to say a few more words about our research program. One obvious goal is to shed light on the subject of contact structures in 3-manifolds using our new methods. Establishing the foundations of our theory was difficult. If we improve and refine our theory enough, we may be able to offer a different approach to the subject of contact structures, sufficiently powerful to give alternative proofs of many known results, and also to uncover new results. Our other goal is to use tight contact structures and tight pure contaminations to prove results about 3-manifolds. To do this it would be best to be able to work entirely with branched surfaces, since branched surfaces are finite combinatorial objects closer to the traditional tools in 3-manifold topology. For the neatest possible theory, it would therefore be best to prove theorems of the following form: (1) If $B \hookrightarrow M$ has properties X, then any pure contamination ξ carried by B is tight; (2) If ξ is a tight pure contamination in M, it is carried by a branched surface B with properties X. Results of this kind were proved for incompressible surfaces and essential laminations, see [6] and [7]. For tight pure contaminations this promises to be more difficult.

2   Eliminating maximal integral surfaces

As a first step towards proving Theorem 1.4, we will show in this section, given a σ-confoliation ζ_0 on a closed 3-manifold M, how to construct a σ-deformation ζ_s from ζ_0 to a contact structure ζ_1.

The elimination of the maximal integral surface will be achieved using a sequence of deformations, each of which either eliminates a disc component from the maximal integral surface F, or removes a 1-handle in F.

We must establish some further notation and definitions before describing precisely the deformations we need. Let $\xi_t : t \in [a, b]$ be a σ-deformation, and $h_t : t \in [a, b]$ a 1-parameter family of diffeomorphisms of the underlying 3-manifold M as in Definition 1.3. We will say that the family h_t puts ξ_t into a *canonical form* (with respect to maximal integral surfaces).

Let $I_1, I_2, \ldots, I_k$ be the sequence of subintervals as in Definition 1.3, corresponding to families ξ_t and h_t as above. For $1 \leq j \leq k - 1$, denote by t_j the common endpoint of the subintervals I_j and I_{j+1}. Points $t_1, t_2, \ldots, t_{k-1}$ will be called *discontinuity points* of a σ-deformation ξ_t.

A σ-deformation with a single discontinuity point t_1 is a *reduction*, if the (strict) inclusion $\Sigma_2 \subset \Sigma_1$ holds; it is an *expansion* in the opposite case. Clearly, each σ-deformation can be decomposed into a combination of reductions and expansions. In this section we describe examples of σ-deformations that we call *controlled reductions*. We use combinations of controlled reductions to show that each σ-confoliation can be connected by a σ-deformation to a contact

structure. We postpone the discussion of expansions until later in this section.

Controlled reductions will be defined in terms of certain charts, which we now define.

Definition 2.1. Let (r, θ, z) be the cylindrical coordinates in $\mathbb{R}^3$, and let $\mathcal{C} = \{(r, \theta, z) \in \mathbb{R}^3 : r < 1, -1 < z < 1\}$. A $\mathcal{C}$-*chart or cylinder chart* for a confoliation ξ in a 3-manifold M is a smooth embedding $\psi : \mathcal{C} \to M$, such that:

(1) the images in M of all curves $\{r = \text{const}, \theta = \text{const}\}$ in $\mathcal{C}$ are transverse to ξ;

(2) the images in M of all curves $\{\theta = \text{const}, z = \text{const}\}$ in $\mathcal{C}$ are everywhere tangent to ξ.

Despite the restrictions imposed by conditions (1) and (2) above, there is quite some freedom in constructing $\mathcal{C}$-charts for a confoliation ξ. We state without proof the following easy fact.

Fact 2.2. *For any disc $D^2 \subset M$ radially foliated by curves tangent to ξ, and for any vector field X transverse to ξ in an open neighborhood U of D^2 and transverse to D^2, there exists a $\mathcal{C}$-chart $\psi : \mathcal{C} \to M$ such that*

(1) $\psi(\{(r, \theta, z) \in \mathcal{C} : z = 0\}) = int(D^2)$;

(2) $\psi(\mathcal{C}) \subset U$;

(3) the image in M of any curve $\{r = const, \theta = const\}$ in $\mathcal{C}$ is an integral curve of the vector field X.

Definition 2.3. Let (x, y, z) be the cartesian coordinates in $\mathbb{R}^3$, and let $\mathcal{B} = \{(x, y, z) \in \mathbb{R}^3 : x, y, z \in (-1, 1)\}$. A $\mathcal{B}$-*chart or box chart* for a confoliation ξ in a 3-manifold M is a smooth embedding $\psi : \mathcal{B} \to M$, such that

(1) the images in M of all curves $\{x = \text{const}, y = \text{const}\}$ in $\mathcal{B}$ are transverse to ξ;

(2) the images in M of all curves $\{x = \text{const}, z = \text{const}\}$ in $\mathcal{B}$ are everywhere tangent to ξ.

The freedom for constructing $\mathcal{B}$-charts for a confoliation ξ is similar to that for $\mathcal{C}$-charts, see Fact 2.2 . We omit the details.

The advantage of working with $\mathcal{C}$-charts and $\mathcal{B}$-charts for a confoliation ξ is that in the coordinates provided by these charts ξ can be expressed in a unique way as the kernel of a 1-form $dz + f(r, \theta, z)d\theta$ ($dz - f(x, y, z)dx$ respectively), for some smooth function f. We will call the function f as above a *slope function*. Deformations of the slope function f then yield deformations of the corresponding confoliation ξ inside a chart neighborhood in M. To perform such deformations we need the following two lemmas.

Lemma 2.4. *Let $\omega = dz - f(x, y, z)dx$ be a 1-form and $\xi = \ker(\omega)$ be the induced plane field in $\mathcal{B}$. Then ξ is a positive confoliation if and only if $\partial f / \partial y(x, y, z) \geq 0$ for each $(x, y, z) \in \mathcal{B}$. Moreover, ξ is contact at a point $(x, y, z) \in \mathcal{B}$ if and only if $\partial f / \partial y(x, y, z) > 0$.*

Lemma 2.5. *Let $\omega = dz + f(r, \theta, z)d\theta$ be a 1-form and $\xi = \ker(\omega)$ be the induced plane field in $\mathcal{C}$.*

(a) The form ω is well-defined and smooth in $\mathcal{C}$ if and only if $f(r, \theta, z) = r^2 \cdot h(r, \theta, z)$ for some smooth function $h : \mathcal{C} \to \mathbb{R}$.

(b) The plane field ξ is a positive confoliation if and only if $\partial f / \partial r(r, \theta, z) \geq 0$ for all points in $\mathcal{C}$ with $r > 0$.

(c) The plane field ξ is a positive contact structure at a point $(r, \theta, z) \in \mathcal{C}$ with $r > 0$ if and only if $\partial f / \partial r(r, \theta, z) > 0$.

(d) The plane field ξ is a positive contact structure at a point $(0, \theta, z) \in \mathcal{C}$ if and only if the function h as in (a) satisfies the condition $h(0, \theta, z) > 0$.

The proof of Lemma 2.4 is straightforward, and can be found in [ET]. We will include the slightly more difficult proof of Lemma 2.5.

Proof. To prove (a), consider coordinates (x, y, z) in $\mathcal{C}$ with $x = r \cos \theta$, $y = r \sin \theta$. In these coordinates we have

$$d\theta = \frac{x\, dy - y\, dx}{x^2 + y^2} \text{ and hence } \omega = dz + \frac{f(x, y, z)}{x^2 + y^2}(x\, dy - y\, dx).$$

For a function $h = f/r^2$ we then have $\omega = dz + h(x, y, z)(x\, dy - y\, dx)$, and therefore ω is smooth if and only if the function $h : \mathcal{B} \to \mathbb{R}$ is well-defined and smooth.

Having proved (a), the other parts of the lemma follow by direct calculation. $\square$

The following lemma establishes the existence of a σ-deformation which eliminates a disc from the maximal integral surface of a σ-confoliation. Any deformation of the kind described in the lemma is called a *controlled reduction* of a disc component.

Lemma 2.6. *Suppose ξ is a σ-confoliation with maximal integral surface F, and suppose D is a disc component of F. There is a σ-deformation of ξ, with one discontinuity point, which reduces the maximal integral surface from F to $F - D$. The deformation can be chosen to be supported in an arbitrarily small neighborhood U of the component D, i.e. it can be chosen to be constant outside of U.*

Proof. Given a neighborhood U of D and a vector field X in U transverse to ξ, whose role is purely technical, it is not difficult to find, essentially using Fact 2.2, a $\mathcal{C}$-chart $\psi : \mathcal{C} \to M$ for ξ, with the following properties.

Properties C:

(1) $\psi(\mathcal{C})$ is contained in U and disjoint from $F - D$;

(2) D is contained in $\psi(\mathcal{C})$ and $\psi^{-1}(D) = \{(r, \theta, z) \in \mathcal{C} : z = 0, r \leq 1 - \epsilon\}$, where ϵ is a small positive number;

(3) the images in M of the curves $\{(r, \theta, z) \in \mathcal{C} : r = \text{const}, \theta = \text{const}\}$ are all contained in the integral curves of the vector field X.

Let $f : \mathcal{C} \to \mathbb{R}$ be the slope function (as defined just before Lemma 2.4) for the σ-confoliation ξ, in some $\mathcal{C}$-chart coordinates as above. Then there exists a smooth family $f_t : t \in [0, 1]$ of functions $f_t : \mathcal{C} \to \mathbb{R}$ having the following properties.

Properties PC:

(1) $f_t = f$ for all $t \leq 1/2$;

(2) for each $t \in [0, 1]$ we have $f_t = h_t \cdot r^2$, with $h_t : \mathcal{C} \to \mathbb{R}$ smooth and smoothly depending on t;

(3) $h_t(0, \theta, z) > 0$ for all points $(0, \theta, z)$ in $\mathcal{C}$ and for all $t > 1/2$;

(4) $\partial f_t / \partial r(r, \theta, z) > 0$ for all $(r, \theta, z) \in \mathcal{C}$ with $r > 0$ and for all $t > 1/2$;

(5) for each $t \in [0, 1]$ we have $f_t = f$ in a neighborhood of the boundary of $\mathcal{C}$ in $\mathbb{R}^3$.

Note that, by Lemma 2.5, conditions (2), (3) and (4) are already satisfied by the slope function f, except that (3) and (4) do not hold for $z = 0$ and $r \leq 1 - \epsilon$ (i.e. at points of the preimage $\psi^{-1}(D)$), where the corresponding inequalities are replaced by equalities. The existence of a family f_t as above is then fairly obvious.

By the properties of functions f_t, it follows from Lemma 2.5 that a 1-parameter family $\xi_t : t \in [0, 1]$ of plane fields in M defined by

$$\xi_t = \begin{cases} \ker(dz + f_t d\theta) & \text{in } \mathcal{C}\text{-chart coordinates inside } \psi(\mathcal{C}) \\ \xi & \text{outside } \psi(\mathcal{C}) \end{cases}$$

is a σ-deformation as required, with the discontinuity point $t_1 = 1/2$. □

The next lemma, analogous to the previous lemma, establishes the existence of a σ-deformation which eliminates a 1-handle or "bridge" from the maximal integral surface F of a σ-confoliation. Thus, a bridge is a rectangular subsurface of F with two opposite sides contained in ∂F and the remaining sides contained in $\text{int}(F)$, except their endpoints. Any deformation of the kind described in the lemma is called a *controlled reduction* of a bridge.

Lemma 2.7. *Suppose ξ is a σ-confoliation with maximal integral surface F, and suppose B is a bridge in F. There is a σ-deformation of ξ, with one discontinuity point, which reduces the maximal integral surface from F to $cl(F - B)$. The deformation can be chosen to be supported in an arbitrarily small neighborhood U of the bridge B, i.e. it can be chosen to be constant outside of U.*

Proof. Since the proof of this lemma is in many respects similar to that of Lemma 2.6, our arguments here are more sketchy.

Let $U \subset M$ be a neighborhood of the bridge B and X a vector field in U transverse to ξ, whose role is again purely technical. Then there exists a $\mathcal{B}$-chart $\psi : \mathcal{B} \to M$ for ξ with the following properties.

Properties B:

(1) $\psi(\mathcal{B})$ is contained in U;

(2) B is contained in $\psi(\mathcal{B})$ and $\psi^{-1}(B) = \{(x, y, z) \in \mathcal{B} : z = 0, |x| \leq 1 - \epsilon, \ \phi_1(x) \leq y \leq \phi_2(x)\}$, where $\epsilon > 0$ is a small number and $\phi_1, \phi_2 : [-1 + \epsilon, \ 1 - \epsilon] \to (-1, 1)$ are two functions with $\phi_1 < \phi_2$; ,

(3) $\psi^{-1}(F) = \psi^{-1}(B) \cup \{(x, y, z) \in \mathcal{B} : z = 0, \ |x| \geq 1 - \epsilon\}$;

(4) the images in M of the curves $\{(x, y, z) \in \mathcal{B} : x = \mathrm{const}, y = \mathrm{const}\}$ are all contained in the integral curves of the vector field X.

See Figure 3.
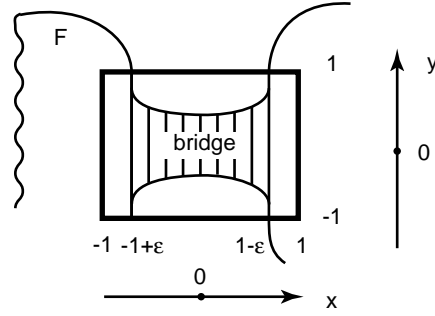

Figure 3: $\mathcal{B}$-chart for a bridge.

Let $f : \mathcal{B} \to \mathbb{R}$ be the slope function for the σ-confoliation ξ, in some $\mathcal{B}$-chart coordinates as above. Then there exists a smooth family $f_t : t \in [0, 1]$ of functions $f_t : \mathcal{B} \to \mathbb{R}$ having the following properties.

Properties PB:

(1) $f_t = f$ for all $t \leq 1/2$;

(2) if $t > 1/2$, then $f_t(x, y, z) = 0$ for all points $(x, y, z) \in \mathcal{B}$ with $z = 0$ and $|x| \geq 1 - \epsilon$ and $\partial f_t / \partial y(x, y, z) > 0$ for all other points in $\mathcal{B}$;

(3) for each $t \in [0, 1]$ we have $f_t = f$ in a neighborhood of the boundary of $\mathcal{B}$ in $\mathbb{R}^3$.

In view of above properties of functions f_t, it follows from Lemma 2.4 that a 1-parameter family $\xi_t : t \in [0, 1]$ of plane fields in M defined by

$$\xi_t = \begin{cases} \ker(dz - f_t dx) & \text{in } \mathcal{B}\text{-chart coordinates inside } \psi(\mathcal{B}) \\ \xi & \text{outside } \psi(\mathcal{B}) \end{cases}$$

is a σ-deformation as required, with the discontinuity point $t_1 = 1/2$. $\qquad\square$

14

Combining simple σ-deformations, it is possible to prove the following:

Proposition 2.8. *For any σ-confoliation ξ there is a σ-deformation ξ_t, $0 \leq t \leq 1$, a composition of finitely many controlled reductions, changing $\xi_0 = \xi$ to a contact structure ξ_1.*

3 Proof of σ-Deformation Theorem

We shall prove Theorem 1.4 by constructing a 2-parameter family of σ-confoliations. Let $\xi_t : t \in [0,1]$ be a σ-deformation between contact structures ξ_0 and ξ_1, as in the theorem. We will use ideas developed in Section 2 to construct a smooth 2-parameter family $\xi_{s,t} : s \in [0,1], t \in [0,1]$ of confoliations with the following properties:

(1) $\xi_{0,t} = \xi_t$ for each $t \in [0,1]$;

(2) $\xi_{s,0}$ and $\xi_{s,1}$ are contact structures for all $s \in [0,1]$;

(3) $\xi_{1,t}$ is a contact structure for each $t \in [0,1]$.

The following proposition yields a 2-parameter family in a slightly more general setting.

Proposition 3.1. *Let $\xi_t : t \in [0,1]$ be a σ-deformation. Then there exists a smooth 2-parameter family $\xi_{s,t}$ of positive confoliations satisfying the following conditions:*

(1) $\xi_{0,t} = \xi_t$ for each $t \in [0,1]$;

(2) both $\xi_{s,0} : s \in [0,1]$ and $\xi_{s,1} : s \in [0,1]$ are combinations of controlled reductions, deforming ξ_0 and ξ_1 respectively into contact structures;

(3) $\xi_{1,t}$ is a contact structure for each $t \in [0,1]$.

Proof of Theorem 1.4 using Proposition 3.1. When ξ_t is a σ-deformation from a contact structure to another contact structure, the 2-parameter smooth family $\xi_{s,t}$ provides a smooth 1-parameter family of contact structures from $\xi_0 = \xi_{00}$ to $\xi_1 = \xi_{10}$ by combining the following 1-parameter families in order, following three consecutive sides of the parameter space $I \times I$ for (s,t) : ξ_{s0}, $0 \leq s \leq 1$ followed by ξ_{1t}, $0 \leq t \leq 1$ followed in turn by $\xi_{(1-s)1}$, $0 \leq s \leq 1$. Now using Gray's Theorem we can replace the smooth 1-parameter family by an isotopy, i.e., a 1-parameter family $\phi_t : M \to M$ of diffeomorphisms such that for all t, $\phi_{t*}(\xi_0) = \xi_t$. $\qquad\square$

From the theorem and proposition, it is now clear that for every σ-confoliation ξ_0, if one gets ξ_1 from ξ_0 via *any* σ-deformation ξ_t, $0 \leq t \leq 1$, then one always obtains the same contact structure up to isotopy, and we call this the *contact structure associated to* ξ_0. Furthermore, this contact structure is actually associated to an equivalence class of σ-confoliations, the equivalence relation being σ-deformation. The existence of the σ-deformation ξ_t from ξ_0 to some contact structure ξ_1 is guaranteed by Proposition 2.8.

We will prove the proposition by proving a sequence of lemmas, some of which are special cases of the proposition. The first lemma will deal with deformations ξ_t having no discontinuity point. The proof in this case will consist of constructing a 2-parameter family of combinations of controlled reductions. Two more lemmas prove the proposition for deformations ξ_t with one discontinuity point. This will require a more skilled use of special charts, to construct smooth parametric families of controlled reductions extended by deformations with no discontinuity points. The final lemma is used as an inductive step for combining 2-parameter families. This will require a delicate argument for showing how to glue to each other 2-parameter families of confoliations corresponding to deformations with single discontinuity points.

Lemma 3.2. *Proposition 3.1 is true if the σ-deformation ξ_t has no discontinuities.*

Proof. By putting such a deformation into a canonical form (as explained at the beginning of Section 2), we will assume that all the confoliations ξ_t have the same maximal integral surface, which we denote by Σ.

Let D be a connected component of Σ which is a disc, and X be a vector field in M transverse to D. Choose a neighborhood U of D such that for each $t \in [0, 1]$ the field X is transverse to ξ_t in U. Then there exists a smooth family $\psi_t : t \in [0, 1]$ of maps $\psi_t : \mathcal{C} \to M$ such that for each $t \in [0, 1]$

(1) ψ_t is a $\mathcal{C}$-chart for ξ_t;

(2) ψ_t has Properties C, as required for a $\mathcal{C}$-chart ψ in the proof of Lemma 2.6.

For each $t \in [0, 1]$, let $f^t : \mathcal{C} \to \mathbb{R}$ be the slope function for ξ_t with respect to a $\mathcal{C}$-chart ψ_t as above. Clearly, the family $f^t : t \in [0, 1]$ of functions depends smoothly on t. Consider then a smooth 2-parameter family $f_s^t : s \in [0, 1], t \in [0, 1]$ of functions $f_s^t : \mathcal{C} \to \mathbb{R}$ with the following properties:

(1) $f_0^t = f_t$ for each $t \in [0, 1]$;

(2) for any fixed $t \in [0, 1]$ the family $f_s^t : s \in [0, 1]$ has Properties PC, as required in the proof of Lemma 2.6 (where the role of parameter t in f_t is taken by s in the family f_s^t).

The existence of a 2-parameter family f_s^t as above is as obvious as the existence of a family f_t in the proof of Lemma 2.6.

Consider a 2-parameter family ξ_s^t of σ-confoliations defined by

$$\xi_s^t = \begin{cases} \ker(dz + f_s^t d\theta) & \text{in cylindrical coordinates induced by } \psi_t \text{ inside } \psi_t(\mathcal{C}) \\ \xi & \text{outside } \psi_t(\mathcal{C}) \end{cases}$$

We shall call any family ξ_s^t obtained by the construction as above a *parametric controlled reduction* of a disc component.

Clearly one can, in a completely analogous way, define the notion of a parametric controlled reduction of a bridge in the maximal integral surface Σ.

The proof is then completed by combining consecutively in s-time a sequence of parametric controlled reductions of discs and/or bridges to obtain the family $\xi_{s,t}$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ $\square$

The following lemma will be used to show that Proposition 3.1 is true for a σ-deformation ξ_t with one discontinuity.

Lemma 3.3. *Let $\xi_t : t \in [0,1]$ be a σ-deformation with exactly one discontinuity point. Suppose the maximal integral surface F_t of ξ_t reduces from Σ_1 to Σ_2, i.e. $\Sigma_2 \subset \Sigma_1$ and $F_t = \Sigma_1$ for $t \leq t_1$ while $F_t = \Sigma_2$ for $t > t_1$ (for some $t_1 \in (0,1)$). Then there exists a smooth 2-parameter family $\xi_{s,t}$ of σ-confoliations such that*

(1) $\xi_{0,t} = \xi_t$ for each $t \in [0,1]$;

(2) Σ_2 is the maximal integral surface for all confoliations $\xi_{1,t}$ with $t \in [0,1]$ and $\xi_{s,1}$ with $s \in [0,1]$;

(3) $\xi_{s,0} : s \in [0,1]$ is a sequence of controlled reductions.

Proof. There is a sequence of operations of deleting a bridge or a disc component, which converts Σ_1 into Σ_2. To prove the lemma, it is then sufficient to consider the case when Σ_2 is obtained from Σ_1 using sequence consisting of a single operation, deleting a bridge or disc component, since the general case will then follow by induction. Therefore, we assume that, to get Σ_2, it is sufficient to delete from Σ_1 a disc component D. The reasoning in the case when Σ_2 is obtained from Σ_1 by deleting a bridge is analogous, and we omit it.

Since D is a connected component of the maximal integral surfaces in ξ_t for all $t \leq t_1$, we may consider a smooth family $\psi_t : t \in [0, t_1]$ of maps $\psi_t : \mathcal{C} \to M$, with the same properties as a family ψ_t of maps considered in the proof of Lemma 3.2. This family may then be extended smoothly to a family $\psi_t : t \in [0, t_1 + \epsilon]$ of maps, for some small positive ϵ, which all satisfy the following two properties:

(1) ψ_t is a $\mathcal{C}$-chart for ξ_t;

(2) $\psi_t(\mathcal{C})$ is disjoint from the surface $\Sigma_2 = \Sigma_1 - D$ for each $t \in [0, t_1 + \epsilon]$.

Since for $t > t_1$ the disc D is no longer an integral surface in ξ_t, we do not require that, for $t > t_1$, D be the core of the each $\mathcal{C}$-charts ψ_t. In fact, for these parameters we do not even require that D be contained in $\psi_t(\mathcal{C})$.

For each $t \in [0, t_1 + \epsilon]$, let $f_t : \mathcal{C} \to \mathbb{R}$ be the slope function for ξ_t with respect to a $\mathcal{C}$-chart ψ_t as above. Consider then a smooth 2-parameter family $f_{s,t} : t \in [0, t_1 + \epsilon], s \in [0, 1]$ of functions $f_{s,t} : \mathcal{C} \to \mathbb{R}$ with the following properties:

(1) $f_{0,t} = f_t$ for each $t \in [0, t_1 + \epsilon]$;

(2) for any fixed $t \in [0, t_1 + \epsilon]$ the family $f_{s,t} : s \in [0,1]$ has Properties PC, as required for a family f_t of functions in the proof of Lemma 2.6 (where the role of parameter t in f_t is taken by s in the family $f_{s,t}$);

(3) for each $t \in [t_1 + \epsilon/2, t_1 + \epsilon]$ and for each $s \in [0, 1]$ we have $f_{s,t} = f_t$.

As before, we regard the existence of such a family $f_{s,t}$ as fairly obvious. In view of Lemma 2.5, it follows from the above properties (1)-(3) of a family $f_{s,t}$ that the 2-parameter family $\xi_{s,t}$ of plane fields in M defined by

$$
\xi_{s,t} = \begin{cases} \ker(dz + f_{s,t}d\theta) & \text{in cylindrical coordinates induced} \\ & \text{by } \psi_t \text{ inside } \psi_t(\mathcal{C}), \text{ if } t \leq t_1 + \epsilon \\ \xi_t & \text{otherwise,} \end{cases}
$$

is well-defined, consists of σ-confoliations, and satisfies the properties (1)-(3) asserted in the lemma. $\qquad\square$

Corollary 3.4. *Proposition 3.1 is true for a σ-deformation ξ_t with exactly one discontinuity.*

Proof. By putting the deformation ξ_t into a canonical form, and by reversing parameter t if necessary, we assume that the maximal integral surface F_t of ξ_t reduces from Σ_1 to Σ_2, i.e. $\Sigma_2 \subset \Sigma_1$ and $F_t = \Sigma_1$ for $t \leq t_1$ while $F_t = \Sigma_2$ for $t > t_1$ (for some $t_1 \in (0, 1)$). We apply the previous lemma to the deformation ξ_t which yields a 2-parameter family $\xi_{s,t}$ with the property that $\xi_{1,t}$ has no discontinuities. We follow the 2-parameter family $\xi_{s,t}$ in s-time by a 2-parameter family as in Lemma 3.2 to eliminate Σ_2. $\qquad\square$

Now we deal with the case of two discontinuities, which contains all the ideas to complete the proof of Proposition 3.1.

Lemma 3.5. *Proposition 3.1 is true for a σ-deformation ξ_t with exactly two discontinuities.*

Proof. Suppose the discontinuities t_1, t_2 are contained in the intervals $(0, 1/2)$ and $(1/2, 1)$ respectively, which we can always assume by reparametrizing in t if necessary. Let $\xi_t^1 : t \in [0, 1/2]$ be defined by $\xi_t^1 := \xi_t$, and similarly $\xi_t^2 : t \in [1/2, 1]$ by $\xi_t^2 := \xi_t$. Then ξ_t^1 and ξ_t^2 are σ-deformations with one discontinuity each. Let $\xi_{t,s}^1 : t \in [0, 1/2], s \in [0, 1]$ and $\xi_{t,s}^2 : t \in [1/2, 1], s \in [0, 1]$ be smooth 2-parameter families constructed as in Lemma 3.3.

The obvious problem with simply glueing families $\xi_{s,t}^1$ and $\xi_{s,t}^2$ into a single family $\xi_{s,t}$ is that deformations $\xi_{s,1/2}^1 : s \in [0, 1]$ and $\xi_{s,1/2}^2 : s \in [0, 1]$ need not to coincide. Both of these deformations are sequences of controlled reductions, but the sequences of bridges and disc components deleted from the common initial maximal integral surface may be very different in the two deformations. To overcome this difficulty, we need to prove the following claim, which provides a link between deformations $\xi_{s,1/2}^1$ and $\xi_{s,1/2}^2$.

Claim. *If all σ-confoliations in a family $\xi_{s,1/2}^2 : s \in [0, 1]$ are sufficiently close to the σ-confoliation $\xi_{0,1/2}^2$, then there exists a smooth 2-parameter family $\eta_{u,v} : u \in [0, 1], v \in [0, 1]$ of positive confoliations, such that:*

(1) $\eta_{u,0} = \xi^1_{u,1/2}$ *for each* $u \in [0,1]$;

(2) $\eta_{0,v} = \xi^2_{v,1/2}$ *for each* $v \in [0,1]$;

(3) all confoliations $\eta_{u,1} : u \in [0,1]$ *and* $\eta_{1,v} : v \in [0,1]$ *are contact structures.*

Before proving the claim, we will finish the proof of the lemma, assuming that the claim is true. To do this, note that a 2-parameter family $\xi_{s,t}$ obtained from Lemma 3.3 can be always constructed in such a way that the restricted deformation $\xi_{s,0} : s \in [0,1]$ consists of confoliations as close to $\xi_{0,0}$ as necessary. This is due to the fact that a controlled reduction of a given bridge or of a given disc component can be made to consist of confoliations arbitrarily close to the initial one. In particular, it is always possible to choose a family $\xi^2_{s,t} : s \in [0,1]$, $t \in [1/2,1]$ so that the restricted deformation $\xi^2_{s,1/2} : s \in [0,1]$ satisfies the assumption of the Claim. The lemma then follows by glueing families $\xi^1_{s,t}$, $\xi^2_{s,t}$ and $\eta_{u,v}$, and performing some cosmetic adjustment in order to make the resulting family $\xi_{s,t}$ smooth, see Figure 3. We omit further details.
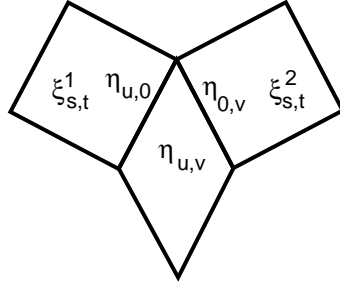

Figure 4: Identifications for domain of 2-parameter family.

Proof of Claim: Put both σ-deformations $\xi^1_{u,1/2} : u \in [0,1]$ and $\xi^2_{v,1/2} : v \in [0,1]$ into canonical forms, with sequences $\Sigma_1, \Sigma_2, \ldots, \Sigma_k = \emptyset$ and $\Sigma'_1, \Sigma'_2, \ldots, \Sigma'_l = \emptyset$ of maximal integral surfaces respectively, so that still $\xi^1_{0,1/2} = \xi^2_{0,1/2}$ and hence also $\Sigma_1 = \Sigma'_1$. Moreover, let $t_1, t_2, \ldots, t_{k-1}$ be the sequence of discontinuity points for $\xi^1_{u,1/2} : u \in [0,1]$. The proof will essentially go by induction with respect to k.

As a first step in this induction we shall construct a smooth 2-parameter family $\eta^1_{u,v} : u \in [0, t_1 + \delta], v \in [0,1]$ of confoliations, where $\delta > 0$ is such that $t_1 + \delta < t_2$, satisfying conditions:

(1) $\eta^1_{u,0} = \xi^1_{u,1/2}$ for each $u \in [0, t_1 + \delta]$;

(2) $\eta^1_{0,v} = \xi^2_{v,1/2}$ for each $v \in [0,1]$;

(3) all confoliations $\eta^1_{u,1} : u \in [0, t_1 + \delta]$ are contact structures;

(4) all confoliations $\eta^1_{t_1+\delta,v} : v \in [0,1]$ are contact at all points of the set $M - \Sigma_2$.

Since, by the assumption, $\xi^1_{u,1/2} : u \in [0,1]$ is a combination of controlled reductions, it follows that Σ_2 is obtained from Σ_1 by deleting a bridge or a

disc component. Assume the first of those two cases, so that $\Sigma_2 = cl(\Sigma_1 - B)$, where B is a bridge in Σ_1. The argument for the second case is analogous, and we will omit it.

Let $\psi : \mathcal{B} \to M$ be the $\mathcal{B}$-chart used for determining the controlled reduction $\xi^1_{u,1/2} : u \in [0, t_1 + \delta]$, as in Lemma 2.7, and let $f_u : u \in [0, t_1 + \delta]$ be the corresponding family of slope functions. Further, let X be the vector field in a neighborhood of the bridge B, corresponding to ψ in the same lemma, Lemma 2.7. If confoliations $\xi^2_{v,1/2} : v \in [0, 1]$ are close enough to the confoliation $\xi^2_{0,1/2}$, then X is transverse to all of them (in the same neighborhood of B), and there exists a smooth 1-parameter family $\psi_v : v \in [0, 1]$ of maps $\psi_v : \mathcal{B} \to M$, with $\psi_0 = \psi$, such that for each $v \in [0, 1]$:

(1) B is contained in $\psi_v(\mathcal{B})$ and

$$\psi_v^{-1}(B) = \{(x, y, z) \in \mathcal{B} : z = g_v(x, y), |x| \leq 1 - \epsilon, \phi_1(x) \leq y \leq \phi_2(x)\},$$

where ϵ, ϕ_1 and ϕ_2 are the same as for ψ, and $g_v : v \in [0, 1]$ is a family of smooth functions $g_v : (-1, 1)^2 \to (-1, 1)$, smoothly depending on v;

(2) $\psi_v^{-1}(\Sigma_2) = \{(x, y, z) \in \mathcal{B} : z = g_v(x, y), |x| \geq 1 - \epsilon\}$, where ϵ and g_v are the same as in (2);

(3) the images by ψ_v of the curves $\{(x, y, z) \in \mathcal{B} : x = \text{const}, y = \text{const}\}$ are all contained in the integral curves of the vector field X.

The main difficulty in constructing a family $\psi_v : v \in [0, 1]$ as above, is to make all the $\mathcal{B}$-charts in the family contain B in the image (condition (2)). The difficulty comes from the fact that, in general, B is not integral in all of the confoliations $\xi^2_{v,1/2} : v \in [0, 1]$. This is where we must make use of the assumption of the claim, which clearly allows the construction of a desired family.

The construction of a family $\eta^1_{u,v}$ of confoliations is now based on the same idea used in previous lemmas dealing with special cases of Proposition 3.1. For each $v \in [0, 1]$, let $f^v : \mathcal{B} \to \mathbb{R}$ be the slope function for the confoliation $\xi^2_{v,1/2}$, in the $\mathcal{B}$-chart ψ_v. Then there exists a smooth 2-parameter family $f^v_u : u \in [0, t_1 + \delta], v \in [0, 1]$ of functions $f^v_u : \mathcal{B} \to \mathbb{R}$ such that:

(1) $f^0_u = f_u$ for each $u \in [0, t_1 + \delta]$, where f_u is the family of slope functions for the controlled reduction $\xi^1_{u,1/2} : u \in [0, t_1 + \delta]$, as mentioned above;

(2) $f^v_0 = f^v$ for each $v \in [0, 1]$;

(3) $\partial f^v_u / \partial y \geq 0$ for each $u \in [0, t_1 + \delta]$ and for each $v \in [0, 1]$;

(4) $\partial f^v_{t_1+\delta} / \partial y > 0$ outside the set $\psi_v^{-1}(\Sigma_2)$;

(5) $f^v_u = f^v$ at a neighborhood of the boundary of $\mathcal{B}$ in $\mathbb{R}^3$, for all $u \in [0, t_1 + \delta]$ and all $v \in [0, 1]$.

The existence of such a family f^v_u of functions is fairly clear.

20

In view of the above properties of functions f_u^v, it follows from Lemma 2.4 that a 2-parameter family $\eta_{u,v}^1 : u \in [0, t_1 + \delta], v \in [0,1]$ of confoliations, defined by

$$
\eta_{u,v}^1 = \begin{cases}
\ker(dz - f_u^v \cdot dx) & \text{in cylindrical coordinates induced} \\
& \text{by } \psi_v \text{ inside } \psi_v(\mathcal{B}) \\
\xi_{v,1/2}^2 & \text{outside } \psi_v(\mathcal{B})
\end{cases}
$$

has the required properties.

To finish the proof of the Claim, note that it is always possible first to choose a deformation $\xi_{v,1/2}^2 : v \in [0,1]$, and then a related family $\eta_{u,v}^1$ as above, so that the confoliations $\eta_{t_1+\delta,v}^1 : v \in [0,1]$ are all arbitrarily close to the confoliation $\eta_{t_1+\delta,0}^1 = \xi_{t_1+\delta,1/2}^1$. This observation allows an inductive construction of a family $\eta_{u,v}$ required in the claim, and so the proof is finished. $\qquad\square$

Proof of Proposition 3.1. The proposition can be proved by induction as follows. Suppose it is true for a deformation with k discontinuities. Replacing $\xi_{t,s}^1$ in the previous lemma by the k-discontinuity deformation, the lemma remains true. Therefore, the proposition is true for a deformation with $k+1$ discontinuities. $\qquad\square$

4   Invariance of tightness

This section is devoted to the proof of Theorem 1.6, showing that tightness of a σ-confoliation is invariant under σ-deformation.

In view of Theorem 1.4 and Proposition 2.8, Theorem reduces to the following.

Theorem 4.1. *Let ξ' be a contact structure obtained from a confoliation ξ by a sequence of controlled reductions. Then ξ' is tight iff ξ is tight.*

Indeed, if $\xi_t : t \in [0,1]$ is a σ-deformation, and ξ_0' and ξ_1' are contact structures obtained from ξ_0 and ξ_1 respectively, by means of sequences of controlled reductions (which exist due to Proposition 2.8), then by Theorem 1.4 those contact structures are isotopic. Hence, Theorem 4.1 and the invariance of tightness under isotopy of contact structures imply that ξ_0 is tight iff ξ_1 is tight, which is exactly the assertion of Theorem 1.6. We therefore now devote the remainder of the section to proving Theorem 4.1. To do this, we need several preparatory results.

Lemma 4.2. *Let D be a disc which is a connected component of the maximal integral surface for a σ-confoliation ξ, and consider a controlled reduction of D which turns ξ into a new σ-confoliation ξ'. Denote by p the point of the underlying 3-manifold M with cylindrical coordinates $(0,0,0)$ in the $\mathcal{C}$-chart ψ used to describe the above controlled reduction. Then the restricted σ-confoliations $\xi|_{M-D}$ and $\xi'|_{M-\{p\}}$ are isotopic by a diffeomorphism $g : M - \{p\} \to M - D$ with the following properties:*

(1) g coincides with the identity outside $\psi(\mathcal{C})$;

(2) each curve in M corresponding by ψ to a curve $\{z = const, \theta = const\}$ in $\mathcal{C}$ is mapped by g to itself.

Proof. Recall from Section 2 that in the $\mathcal{C}$-chart ψ used for the description of the controlled reduction of D we have

$$\xi = \ker(dz + f \cdot d\theta) \quad \text{and} \quad \xi' = \ker(dz + f' \cdot d\theta),$$

where the functions $f, f' : \mathcal{C} \to \mathbb{R}$ satisfy the following conditions:

(1) $f(r, \theta, z) = r^2 \cdot h(r, \theta, z)$ and $f'(r, \theta, z) = r^2 \cdot h'(r, \theta, z)$ for some smooth functions $h, h' : \mathcal{C} \to \mathbb{R}$;

(2) $h' > 0$ everywhere in $\mathcal{C}$, and $h > 0$ except at the set $\{z = 0, r \leq 1 - \epsilon\} = \psi^{-1}(D) \subset \mathcal{C}$ at which $h = 0$;

(3) $\partial f'/\partial r(r, \theta, z) > 0$ for all points (r, θ, z) in $\mathcal{C}$ with $r > 0$;

(4) $\partial f/\partial r(r, \theta, z) > 0$ for all points (r, θ, z) in $\mathcal{C} - \psi^{-1}(D)$ with $r > 0$;

(5) $f = f'$ at a neighborhood of the boundary of $\mathcal{C}$ in $\mathbb{R}^3$.

Define a map $u : \mathcal{C} - \{z = 0, r \leq 1 - \epsilon\} \to \mathcal{C} - \{(0,0,0)\}$ of the form $u(r, \theta, z) = (w(r, \theta, z), \theta, z)$, by putting

$$w(r, \theta, z) = (f'_{\theta,z})^{-1} \circ f_{\theta,z}(r),$$

where $f_{\theta,z}(r) = f(r, \theta, z)$ and $f'_{\theta,z}(r) = f'(r, \theta, z)$. Due to the properties (1)-(5) above u is a well-defined homeomorphism equal to the identity at a neighborhood of the boundary of $\mathcal{C}$ in $\mathbb{R}^3$. Moreover, u is easily seen to be regular at all points (r, θ, z) in its domain with $r > 0$. Therefore, to prove that u is a diffeomorphism, it is sufficient to show that it is regular at points with $r = 0$. This follows however from the fact that u can be expressed as a composition of the map $(r, \theta, z) \to (r \cdot h^{1/2}, \theta, z)$ with the inverse of the map $(r, \theta, z) \to (r \cdot (h')^{1/2}, \theta, z)$, and both those maps are regular at points with $r = 0$ and $z \neq 0$ due to property (2).

A straightforward argument shows that, by the definition, the diffeomorphism u maps the structure ξ to the structure ξ'. Since moreover, by property (5), u is equal to the identity close to the boundary of $\mathcal{C}$, it can be extended to a map g as required, by putting $g = id$ outside the chart neighborhood $\psi(\mathcal{C})$. This finishes the proof. $\qquad\square$

Lemma 4.3. *Let $\xi_t : t \in [0,1]$ be a controlled reduction of a bridge B, as in Section 2. Denote by F_0 and F_1 the maximal integral surfaces for ξ_0 and ξ_1 respectively, and by γ an arc in F_0 defined by $\gamma = \psi(\{(x, y, 0) \in B : |x| < 1 - \epsilon, f_1(x, y, 0) = 0\})$, where $\psi : B \to M$ is the B-chart used to describe the controlled reduction ξ_t, and the rest of notation is as in Section 2. Then the contact structures $\xi_0|_{M-F_0}$ and $\xi_1|_{(M-(F_1\cup\gamma)}$ are isotopic by a contactomorphism which coincides with the identity outside the set U supporting the deformation ξ_t.*

The argument for proving this lemma is analogous as that for Lemma 4.2, and we omit it.

Let F be a compact surface, each component of which has nonempty boundary. A *core* in F is a graph Γ smoothly embedded in the interior of F, whose image is a deformation retract of F. In what follows we shall identify Γ with its image in F.

Crucial for proving Theorem 4.1 is the following

Proposition 4.4. *(1) Let ξ be a σ-confoliation in M, F its maximal integral surface, and Γ a core in F. Then, given a neighborhood U of F, there exists a sequence of controlled reductions which changes ξ into a contact structure ξ' in such a way that the restricted contact structures $\xi|_{M-F}$ and $\xi'|_{M-\Gamma}$ are isomorphic by a contactomorphism h equal to the identity outside U.*

(2) A contactomorphism h as above can be chosen to satisfy the following additional property. For arbitrarily small regular neighborhood W of the core Γ in M the restriction of h to the preimage $h^{-1}(M-W)$ extends to a diffeomorphism of M. Regularity of a neighborhood W means that it is the interior of a compact submanifold with boundary in M.

Proof. (1) Note first that, by the definition of a core, the pair (F, Γ) is diffeomorphic to the pair $(N\Gamma, \Gamma)$, where $N\Gamma$ is a small tubular neighborhood in F of Γ. It follows that there exists a decomposition of the surface F into a union of a disjoint family $D_v : v \in V\Gamma$ of discs corresponding to the vertices of Γ, and a disjoint family $B_e : e \in E\Gamma$ of bridges corresponding to the edges of Γ. Such a decomposition may be easily chosen to have an additional property that for each edge e of Γ the intersection $B_e \cap \Gamma$ is a closed subarc contained in the interior of e, and it meets the boundary ∂B_e transversely.

We will construct a required system of controlled reductions using a system of bridges and discs as above. Two steps of this construction will consist of first reducing all the bridges and then all the resulting discs. We describe these two steps separately.

Step 1. Reduction of the bridges B_e.

Choose a system of $\mathcal{B}$-charts $\psi_e : \mathcal{B} \to M$ around the bridges B_e which satisfy Properties B (as in the proof 2.7) and:

(1) the chart neighborhoods $\psi_e(\mathcal{B})$ in M are pairwise disjoint;

(2) for each edge e we have $\psi_e^{-1}(\Gamma) = \{(x, y, z) \in \mathcal{B} : y = 0, z = 0\}$.

For each edge e of Γ consider a controlled reduction of the bridge B_e, determined as in Section 2 with respect to the $\mathcal{B}$-chart ψ_e. Denote by $\xi_t : t \in [0, 1/2]$ a composition of the controlled reductions as above. We may assume for each e the following additional property: if $\xi_{1/2} = \ker(dz - f^e dx)$ in the chart ψ_e, then $f^e(x, 0, 0) = 0$ for each $x \in (-1, 1)$. It follows then from Lemma 4.3 applied inductively that the contact structures $\xi|_{M-F}$ and $\xi_{1/2}|_{M-(\cup D_v \cup \Gamma)}$ are isotopic as required.

Step 2. Reduction of the discs D_v.

Note that the maximal integral surface of the σ-confoliation $\xi_{1/2}$, as obtained in Step 1, is a union of the family $D_v : v \in V\Gamma$ of pairwise disjoint discs. Choose a system of $\mathcal{C}$-charts $\psi_v : \mathcal{C} \to M$ around those discs which satisfy Properties C (as in the proof of Lemma 2.6) and:

(1) the chart neighborhoods $\psi_v(\mathcal{C})$ in M are pairwise disjoint;

(2) for each vertex v there exists a finite set Θ_v such that $\psi_v^{-1}(\Gamma) = \{(r, \theta, z) \in \mathcal{C} : z = 0, \theta \in \Theta_v\}$.

For each vertex v of Γ consider a controlled reduction of the disc D_v determined as in Section 2 with respect to the $\mathcal{C}$-chart ψ_v. Denote by $\xi_t : t \in [1/2, 1]$ a composition of the controlled reductions above. By Lemma 4.2, for each v there is a diffeomorphism $g_v : \psi_v(\mathcal{C}) - D_v \to \psi_v(\mathcal{C}) - \psi_v(0, 0, 0)$ which maps the restricted structure $\xi_{1/2}$ to the restricted structure ξ_1. Since, by the same lemma, the map g_v maps each curve $\psi_v(\{\theta = \text{const}, z = \text{const}\})$ to itself, we may view it as a contactomorphism between the structures $\xi_{1/2}|_{\psi_v(\mathcal{C}) - (D_v \cup \Gamma)}$ and $\xi_1|_{\psi_v(\mathcal{C}) - -}$. Since moreover each of the maps g_v coincides with the identity close to the boundary of $\psi_v(\mathcal{C})$, the system of these maps can be extended by the identity to a contactomorphism g between the structures $\xi_{1/2}|_{M - (\cup D_v \cup \Gamma)}$ and $\xi_1|_{M - \Gamma}$. Together with Step 1, this gives an isotopy between the contact structures $\xi|_{M - F}$ and $\xi_1|_{M - \Gamma}$, and therefore the composition $\xi_t : t \in [0, 1]$ of the deformations from Steps 1 and 2 satisfies the assertions of the proposition.

(2) The second assertion in 4.4 follows from the fact that a diffeomorphism $h : M - F \to M - \Gamma$ is in fact obtained as the final map in a smooth 1-parameter family of embeddings of the complement $M - F$ into M, starting at the inclusion map. This 1-parameter family can be associated to controlled reductions described in the proof of part (1). Restriction of this 1-parameter family to the preimage $h^{-1}(M - W)$ then yields an isotopy of embeddings of a submanifold $h^{-1}(M - W)$ into M, and this can be extended to an isotopy of diffeomorphisms of M by the standard results concerning isotopy extensions. $\square$

Proof of Theorem 4.1. Call any disc Δ as in Definition 1.5 an *overtwisting disc* of a σ-confoliation (compare [4]). Then tightness boils down to nonexistence of an overtwisting disc. To prove Theorem 4.1 we shall show that such a disc exists for ξ iff it exists for ξ'. This follows fairly directly from Proposition 4.4 as follows.

Assume first that there is an overtwisting disc for ξ, and call it Δ. Since the boundary $\partial\Delta$ is disjoint with the maximal integral surface F for ξ, we can choose a neighborhood U of F disjoint from $\partial\Delta$. By Proposition 4.4, it is then possible to deform ξ into a contact structure ξ'' by a sequence of controlled reductions with supports contained in U, so that ξ and ξ'' coincide outside U. Then Δ is clearly an overtwisting disc for ξ''. However, since due to Theorem 1.3 the contact structures ξ'' and ξ' are isotopic, it follows that there exists an overtwisting disc for ξ'.

Conversely, assume that Δ is an overtwisting disc for ξ'. By Theorem 1.3, we may assume that ξ' is obtained from ξ by a sequence of controlled reductions as in Proposition 4.4, with a core Γ and a contactactomorphism h as in the statement of this proposition. Now, note that if the boundary curve $\partial\Delta$ is deformed within the class of smooth closed embedded curves tangent to ξ', then it still bounds an overtwisting disc for ξ'. Moreover, it is always possible to deform $\partial\Delta$ as above so that after deformation it is disjoint from the graph Γ. Therefore we may assume that the boundary $\partial\Delta$ of an overtwisting disc Δ for ξ' is disjoint from Γ.

Let $h : M - F \to M - \Gamma$ be a contactomorphism as prescribed by Proposition 4.4. Then the curve $h^{-1}(\partial\Delta)$ is clearly disjoint with F and tangent to ξ. Choose a regular neighborhood W of the graph Γ disjoint with $\partial\Delta$. Let $H : M \to M$ be a diffeomorphism whose restriction to the preimage $h^{-1}(M - W)$ coincides with h, as prescribed by part (2) of Proposition 4.4. Then the disc $H^{-1}(\Delta)$ is easily seen to be transverse to ξ at its boundary and hence it is an overtwisting disc for ξ. $\square$

5   Contaminations and splitting

In this section we give a precise definition of a contamination as an object with a smooth structure. We also discuss various operations on pure contaminations, especially an operation of splitting. A special case of this operation, applied to contact structures, was described in the introduction.

A *manifold with inward cusps* is a topological oriented 3-manifold V with boundary ∂V and with distinguished collection $\mathcal{C}(V)$ of pairwise disjoint simple closed curves in ∂V called *cusp curves*. Each cusp curve $c \in \mathcal{C}(V)$ is furthermore equipped with a germ g_c of a homeomorphic identification of the two sides of a thin neighborhood N_c of c in ∂V. We assume also that for each $c \in \mathcal{C}(V)$ the germ g_c extends the identity map on c (intuitively, g_c is a germ of a glueing map for V which identifies the two sides of the strip N_c after folding it along c), see Figure 5.
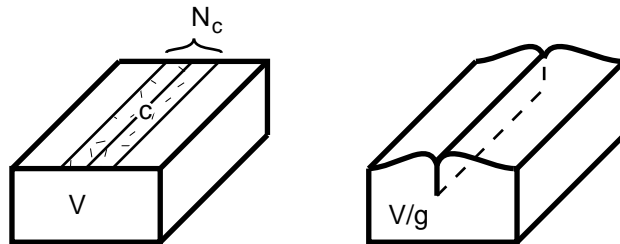

Figure 5: The smooth germ at the cusp locus.

A smooth structure on a manifold with inward cusps V is given by the following data:

(1) ordinary charts (including boundary charts) at open neighborhoods in V disjoint with $\mathcal{C}(V)$;

(2) germs of charts at points of $\mathcal{C}(V)$ described more precisely as follows.

Let $c \in \mathcal{C}(V)$ be a cusp curve of V and let $p \in c$. For a homeomorphism g representing germ g_c let V/g be the quotient of V with respect to g viewed as a glueing map. A *germ of a chart at p* is a germ of a homeomorphism $U_p \to U \subset \mathbb{R}^3$, where U_p is a sufficiently small neighborhood of p in V/g (so that U_p is homeomorphic to an open set in $\mathbb{R}^3$).

We of course assume that transition maps for all charts and germs of charts are smooth. We notice that a smooth structure on V induces smooth structures on the components of the horizontal boundary $\partial_h V$ and on the cusp curves in $\mathcal{C}(V)$. It then turns out that the germs g_c are smooth. On the other hand, the boundary ∂V as a whole does not inherit a canonical smooth structure. Charts for ∂V at points of cusp curves are not canonically determined.

With the above definition of a smooth structure on a manifold with inward cusps V, one can speak of various smooth objects on V, like functions, vector fields, etc. This justifies the definition of a (smooth) contamination in V given in the introduction (Definition 1.8). It also becomes clear what it means for V to be smoothly embedded in a 3-manifold M. Strictly speaking, an embedding $V \hookrightarrow M$ is not injective near cusp curves; it respects identifications at thin neighborhoods of cusp curves induced by the germs g_c. Finally, *diffeomorphisms* of manifolds with inward cusps (as well as general smooth maps between them) and *isomorphisms of contaminations* are now well-defined.

Given a branched surface B, a fibered neighborhood $V(B)$ of B, as described in the introduction, is a well-defined (up to diffeomorphism) smooth manifold with inward cusps. Furthermore, if B is embedded in a 3-manifold M, its fibered neighborhood $V(B)$ is also embedded in M, uniquely up to isotopy of embeddings. Similarly, a contamination ξ carried by a branched surface B embedded in M is well-defined only up to isotopy of embeddings $V(B) \hookrightarrow M$.

A *(pure) deformation of contaminations* on a fixed manifold with inward cusps V is a 1-parameter family ξ_t of (pure) contaminations smoothly depending on the parameter t. It is worth noting that pure contaminations connected by a pure deformation (unlike contact structures on closed manifolds connected by a pure deformation) need not be isomorphic. One of the reasons is that such contaminations may behave in a significantly different way near the boundary (the rate of convergence to tangency when approaching $\partial_h V$ is a sensitive invariant of isomorphism). On the other hand, even the restricted contact structures in $\mathrm{int}(V)$ for contaminations connected by a pure deformation probably need not be isomorphic. For our purposes in this paper it is sufficient to view pure contaminations up to an even weaker relation than pure deformation. To describe this relation, we need some further definitions.

A *σ-contamination* in a manifold with inward cusps V is a smooth plane field ξ in V tangent to $\partial_h V$, for which there exists a compact surface F (with piecewise smooth boundary) smoothly embedded in V, such that

(1) $\mathrm{int}(F)$ is embedded in $\mathrm{int}(V)$, the intersection $F \cap \partial V$ is contained in cusp curves $\mathcal{C}(V)$ and F is tangent to $\partial_h V$ at the intersection;

(2) the intersection $\partial F \cap \mathcal{C}(V)$ is a finite number of arcs and the components of $[\partial F \cup \mathcal{C}(V)] - \mathrm{int}[\partial F \cap \mathcal{C}(V)]$ are smooth closed curves (i.e. the segments of $\partial F \cap \mathrm{int}(V)$ smoothly attach to the segments of $\mathcal{C}(V) - F$);

(3) for each connected component F_i of F, $\partial F_i - \mathcal{C}(V) \neq \emptyset$;

(4) ξ is everywhere tangent to F and it is contact on $\mathrm{int}(V) - F$.

(5) At points where F meets $\mathcal{C}(V)$, $\partial_h V \cup F$ is smooth.

The surface F is called the *maximal integral surface* of a σ-contamination.

Observe that both pure contaminations and σ-confoliations are special cases of σ-contaminations. Condition (3) in the above definition corresponds, in the definition of a σ-confoliation, to the condition that the maximal integral surface of a σ-confoliation cannot have a closed component. Condition (2) allows cutting V along F to get another manifold with inward cusps, as will be described later. Components F_i intersecting $\mathcal{C}(V)$ smoothly attach to components of $\partial_h V$, making the entire locus of integrability of ξ a smooth branched surface, improperly embedded in M.

A σ-*deformation of* σ-*contaminations* is a smooth 1-parameter family $\xi_t : t \in [a, b]$ of plane fields on a manifold with inward cusps V, satisfying the following conditions:

(1) for each t, ξ_t is a σ-contamination on V;

(2) there exists a 1-parameter family $h_t : t \in [a, b]$ of diffeomorphisms of V, a partition of the interval $[a, b]$ into a finite number of nontrivial disjoint consecutive subintervals $I_1, I_2, \ldots, I_k$, and a sequence $\Sigma_1, \Sigma_2, \ldots, \Sigma_k$ of compact surfaces in V such that:

(a) for each $t \in \mathrm{int}(I_j)$ we have $h_t(F_t) = \Sigma_j$ (where F_t is the maximal integral surface of ξ_t);

(b) for each $1 \leq j \leq k - 1$ the surfaces Σ_j and Σ_{j+1} are distinct, and one of them is contained in the other;

(c) if $\Sigma_j \subset \Sigma_{j+1}$ (or conversely $\Sigma_{j+1} \subset \Sigma_j$), then the closure $\mathrm{cl}((\Sigma_{j+1} - \Sigma_j))$ ($\mathrm{cl}(\Sigma_j - \Sigma_{j+1})$ respectively) is a compact surface with piecewise smooth boundary;

(d) each component of the larger of the unions $\Sigma_j \cup \mathcal{C}(V)$ and $\Sigma_{j+1} \cup \mathcal{C}(V)$ is obtained from the smaller one by attaching 1-handles and adding 0-handles, where 1-handles may be attached to $\mathcal{C}(V)$.

Notice that a σ-deformation of σ-confoliations (as defined in the introduction) is a special case of a σ-deformation for σ-contaminations.

In analogy with the results of Section 2 we get the following lemma which shows that it is possible to eliminate the maximal integral surface in any σ-contamination, by means of a σ-deformation.

Lemma 5.1. *For each σ-contamination ξ there exists a σ-deformation which modifies ξ to yield a pure contamination.*

To prove this lemma one can work with special charts as in Section 2, performing controlled reductions associated to those charts. This requires an obvious adaptation of types of special charts to the situations when the integral surface meets the cusp locus $\mathcal{C}(V)$. We omit the details.

By the arguments as in Section 3, generalized as indicated above, one gets also the following extension of the σ-Deformation Theorem (Theorem 1.4).

Theorem 5.2. *Any two pure contaminations which are connected by a σ-deformation of σ-contaminations are also connected by a pure deformation.*

We now turn to describing an operation of splitting for pure contaminations. This operation consists of a sequence of two operations: first, a σ-deformation which introduces a nonempty maximal integral surface followed, second, by cutting on this maximal integral surface. More precisely, the cutting is applied to a given σ-contamination ξ on a manifold with inward cusps V, with maximal integral surface F, and cutting V on F yields a new smooth manifold with inward cusps V'. The cusp locus $\mathcal{C}(V')$ of V' is then equal to $[\partial F \cup \mathcal{C}(V)] -$ int$[\partial F \cap \mathcal{C}(V)]$ and the germs of identification close to new parts of $\mathcal{C}(V')$ are induced by the identification of the two copies of F which reverse the cutting operation. There is also an obvious smooth structure on V' induced from V, and the canonical smooth map $s : V' \to V$. We obtain a plane field ξ' on V' induced by s, which is clearly a pure contamination. We will call ξ' the *pure contamination obtained by cutting ξ* on the maximal integral surface.

Definition 5.3. A *splitting* of a pure contamination ξ in V is an operation consisting of a σ-deformation of ξ to a σ-contamination ξ' with maximal integral surface F, followed by cutting ξ' (and its support V', which is the same as the support V of ξ) to yield a new pure contamination ξ'' with support V''.

If V'' is obtained from V by cutting as above, then an embedding of V in a 3-manifold M induces an embedding of V'' in M, well-defined up to isotopy of embeddings. Thus, any splitting of a pure contamination embedded in M yields a pure contamination embedded in M. One case of interest to us occurs when splitting a contact structure in M yields a pure contamination in M with product complementary pieces, as defined in the introduction.

The next section contains descriptions of certain splittings used to prove Theorem 1.9.

6 Contact structures are carried by branched surfaces

In this section we prove Theorem 1.9, showing that any contact structure ξ defined on a 3-manifold M is σ-deformable to a σ-confoliation, which, when cut

on its maximal integral surface, is a pure contamination carried by a branched surface. Recall that the statement of our theorem is actually stronger: Given any 1-foliation Ω transverse to ξ, we can find a σ-deformation ξ_t transverse to Ω at all times t, $0 \leq t \leq 1$, such that the maximal integral surface of ξ cuts leaves of Ω into intervals.

Proof. Theorem 1.9. We begin by choosing any surface F everywhere transverse to Ω which cuts the leaves of Ω into intervals. Finding such a surface is not difficult. The idea of the proof is to σ-deform ξ to ξ', introducing integral surfaces near F such that for every intersection of a leaf of Ω with F there is an intersection nearby on the leaf of Ω with the integral surface of ξ'.

If the surface F is in a suitable general position with respect to ξ, then the plane field ξ induces on F a singular line field whose integral curves form the *characteristic foliation* or *induced foliation*. In fact, the surface F can be isotoped slightly such that the induced foliation on F is a singular foliation with only elliptic and hyperbolic singularities, all of these being in the interior of F. The induced foliation on F may have tangencies on ∂F, i.e. points where the induced foliation is tangent ot ∂F, and these tangencies are of two types as shown in Figure 6, again assuming a suitable general position.

At the singularities F is tangent to ξ. For every elliptic singularity of $F \hookrightarrow M$ we choose a $\mathcal{C}$-chart centered at the singularity. Recalling $\mathcal{C} = \{(r, \theta, z) \in \mathbb{R}^3 : r < 1, \ -1 < z < 1\}$, we choose a $\mathcal{C}$-chart such that $\{z = 0\}$ is a disc in F, the curves $r = \mathrm{const}$, $\theta = \mathrm{const}$ are segments in leaves of Ω, and radii $\theta = \mathrm{const}$, $z = 0$ are contained in leaves of the induced foliation on F. Here, abusing notation, we identify $\mathcal{C}$ with its image in M. This choice of a $\mathcal{C}$-chart is possible by Fact 2.2.

For every $\mathcal{C}$-chart $\mathcal{C}_i$ associated to an elliptic point on F, we may assume after possibly reparametrizing r, that there is a σ-deformation which introduces an integral disc at $r \leq 1/2$, $z = 1/2$, i.e. to one side of F. Also, we may assume that the support of the σ-deformation is contained in the interior of the $\mathcal{C}$-chart, and is disjoint from $\{z < 1/8\}$.

At each hyperbolic singularity of the induced foliation on F, we choose a disc $D \hookrightarrow M$ transverse to Ω and tangent to F at the singularity such that the induced foliation on D has exactly one elliptic point at the same singularity. After possibly replacing D by a smaller concentric disc, we can then choose a $\mathcal{C}$-chart such that $\{z = 0\} = D$, the curves $r = \mathrm{const}$, $\theta = \mathrm{const}$ are segments in leaves of Ω, and radii $\theta = \mathrm{const}$, $z = 0$ are contained in leaves of the induced foliation on D. Again this is possible by Fact 2.2. We can further require, again by possibly replacing D by a smaller disc, that each segment $r = \mathrm{const}$, $\theta = \mathrm{const}$ intersects each of F and D exactly once, and that $F \cap \mathcal{C} \subset \{|z| < 1/4\}$. By reparametrizing r, we can suppose that there is a σ-deformation with support in $z > 3/8$ introducing an integral surface $r \leq 1/2$, $z = 1/2$.

For every $\mathcal{C}$-chart $\mathcal{C}_i$ associated to an elliptic or hyperbolic point in F, we project the disc $\{z = 1/2, r \leq 1/4\}$ vertically along the 1-foliation to a disc

in F, and we denote by Δ the union of these discs, assumed disjoint, in F. In fact, in the above construction we may arrange that there are exactly four tangencies of the induced foliation with every boundary component of $\partial(F - \mathring{\Delta})$ corresponding to a hyperbolic singularity and no tangencies on any boundary component corresponding to an elliptic singularity. Further, the tangencies are all of the same type; locally near a tangency point, a leaf of the induced foliation tangent to the boundary lies inside $F - \mathring{\Delta}$. We want the induced foliation on $F - \mathring{\Delta}$ to be a foliation by intervals. That is not the case, since the induced foliation might contain closed leaves or non-compact leaves. To achieve our goal we will enlarge Δ by adding more $\mathcal{C}$-charts.

For any $p \in (F - \Delta) \hookrightarrow M$, we can choose a $\mathcal{C}$-chart with $\{p\} = \{r = 0, \ z = 0\}$, such that the disc $D = \{z = 0\}$ becomes a disc transverse to Ω, such that the curves $r = \text{const}$, $\theta = \text{const}$ are segments in leaves of Ω, and such that radii $\theta = \text{const}$, $z = 0$ are contained in leaves of the induced foliation on D. As before, we also require, again by possibly replacing D by a smaller disc, that each segment $r = \text{const}$, $\theta = \text{const}$ intersects each of F and D exactly once, and that $F \cap \mathcal{C} \subset \{|z| < 1/4\}$. By reparametrizing r, we can suppose that there is a σ-deformation with support in $z > 3/8$ introducing an integral surface $r \leq 1/2$, $z = 1/2$.

For a suitable finite collection of $\mathcal{C}$-charts as above, mutually disjoint and also disjoint from $\mathcal{C}$-charts associated to elliptic and hyperbolic singularities on F, if we add to Δ the discs $\{z = 1/2, r \leq 1/4\}$ projected to F, we ensure that the induced foliation of $F - \mathring{\Delta}$ has only interval leaves, with finitely many tangencies at the boundary. We can assume that there are exactly two tangencies (of the same type as for hyperbolic singularities) at boundary components corresponding to discs of Δ for a non-singular point, see Figure 6.

Now, let F_0 be a slightly larger surface than F, containing F in its interior. We may assume that the induced foliation on $F_0 - \mathring{\Delta}$ has the same properties as the one on $F - \mathring{\Delta}$, as described above. Then $F - \mathring{\Delta}$ can be covered by finitely many (rectangular) discs of the form $R_i = \{(x, y) \in \mathbb{R}^2 : x, y \in (-1, 1)\}$, with the property that the segments $y = \text{const}$ lie in leaves of the induced foliation on F_0, and contain entire leaves of the induced foliation on $F - \mathring{\Delta}$, see Figure 6. We also require that for each $p \in F_0$, p belongs to a most two R_i's.

For each R_i we now find a $\mathcal{B}$-chart $\mathcal{B}_i = \{(x, y, z) \in \mathbb{R}^2 : x, y, z \in (-1, 1)$ such that (the images of) the curves $\{x = const, y = const\}$ in $\mathcal{B}_i$ are contained in the leaves of Ω; such that projection of $\mathcal{B}_i$ to F_0 along the leaves of Ω coincides with R_i; such that the (x, y)-coordinates in $\mathcal{B}_i$ are induced from the (x, y)-coordinates in R_i via this projection; and such that the curve $\{y = 0\}$ in R_i coincides with (the image of) the curve $\{z = 0, y = 0\}$ in $\mathcal{B}_i$.

If the rectangles R_i are sufficiently thin in their y-coordinates (which can always be assumed), we can choose the $\mathcal{B}_i$'s so that for each i the rectangle R_i is contained in (the image of) $\mathcal{B}_i$, and for any cylindrical chart $\mathcal{C}_j$ described above we have $\mathcal{B}_i \cap \mathcal{C}_j \subset \{|z| < 1/8\}$, where the z-coordinate is in $\mathcal{C}_j$. This means the box charts are away from the support of the σ-deformations in the
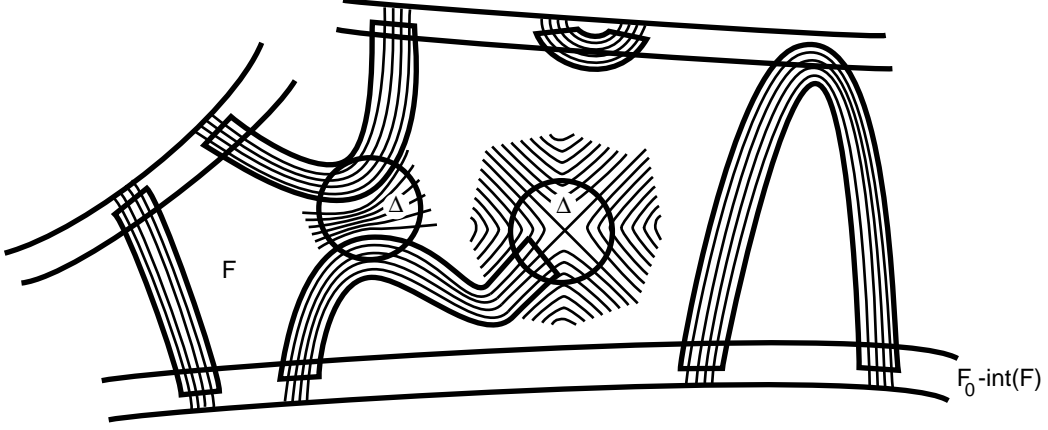
Figure 6: Some rectangles R_i.

cylinder charts.

For every $\mathcal{B}_i$ consider a *central rectangle* P_i in $\mathcal{B}_i$ defined by $P_i = \{z = 0\}$. Observe that due to our assumptions, any rectangle P_i (viewed as a subset of M) intersects F_0 exactly at the curve $\{z = 0, y = 0\}$, and that the two parts of P_i corresponding to $y > 0$ and $y < 0$ lie on opposite sides of the surface F_0, see Figure 7. From this we easily deduce that the central rectangles P_i are pairwise disjoint in M, since the parts that overlap when projected to F_0 lie on opposite sides of F_0.
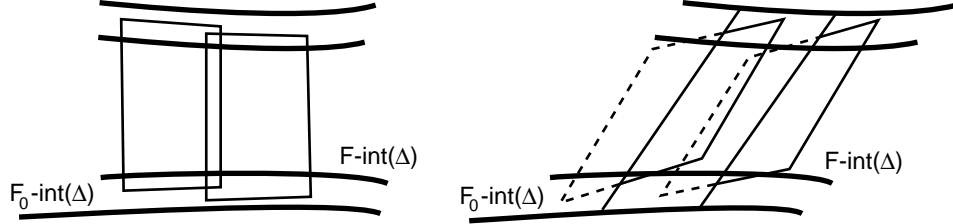
Figure 7: Overlapping rectangles R_i and corresponding central rectangles P_i.

Observe now that, by restricting to a smaller range of the z-coordinate around $z = 0$, we may choose smaller $\mathcal{B}$-charts $\mathcal{B}'_i$ which contain the P_i's as central rectangles, and which are pairwise disjoint. We will show that it is possible to perform a σ-deformation supported in $\mathcal{B}'_i$ which produces a disc component of the maximal integral surface contained in the central rectangle P_i and as close to this rectangle as we wish. Note that this is sufficient for completing the proof, since projections of such discs to F_0 still cover $F - \overset{\circ}{\Delta}$.

Construction of a σ-deformation as required is very similar to the construction of a controlled reduction that eliminates a bridge, as in Section 2. It can be performed by deforming the slope function for the contact structure ξ expressed in the corresponding $\mathcal{B}$-chart $\mathcal{B}'_i$. The differences are that now we introduce a piece of integral surface instead of eliminating it, and that the whole process

31

concerns an isolated disc component of the maximal integral surface instead of a bridge that connects two components of a surface where ξ is integrable. Nevertheles, the details are completely analogous to those described in Section 2, and we omit them. $\square$

7   Three examples in S^3

Given a contact structure ξ in a 3-manifold, if we find a branched surface carrying ξ by the procedure described in Section 6, we will almost certainly produce an unnecessarily complicated branched surface. In practice, one should look for the simplest and most natural branched surface reflecting the geometry of the contact structure. However, we shall see that if the branched surface is not "fine" enough, it may not encode the tightness or overtwistedness of the contact structure(s) carried by it. Clearly, examples of contact structures in S^3 are among the most fundamental in the theory, so we examine, in this section, branched surfaces for some contact structures in S^3. Our explanations will be somewhat informal.

We begin with a model of S^3 suited to our purposes. Namely, we regard S^3 as a quotient space obtained from $T^2 \times [-1, 1]$, where T^2 is a torus, by identifying slope ∞ circles in $T^2 \times \{-1\}$ to points and similarly collapsing slope 0 circles to points in $T^2 \times \{1\}$. In our figures we will represent the sphere as a cube $[-1, 1] \times [-1, 1] \times [-1, 1]$ in $\mathbb{R}^3$, see Figure 8, where the squares $y = \text{const}$ represent the tori.
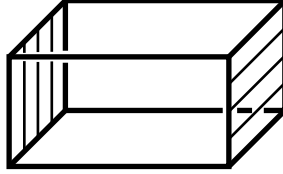

Figure 8: The 3-sphere.

Now we describe three positive contact structures on S^3. All three have horizontal lines $p \times [-1, 1]$ as Legendrians, $p \in T^2$. In our model these appear as lines $z = \text{const}$, $x = \text{const}$. The first contact structure, which we denote ξ_0 is defined by the property that planes rotate $\pi/2$ going from $y = -1$ to $y = +1$ on each horizontal Legendrian. The contact structures ξ_1 and ξ_2 are similar, but with a rotation of $3\pi/2$ and $5\pi/2$, respectively, on horizontal Legendrians. See Figures 9, 10. The reader may recognize these structures as the standard structure ξ_0, the once overtwisted structure ξ_1, and the twice overtwisted structure ξ_2. We note that adding additional twists does not give new overtwisted structures, though this may not be obvious to the reader.

We also remark that the Hopf fibration can easily be recognized in our model of S^3. The circle leaves are given by lines of slope 1 in each plane $y = \text{const}$, namely lines of the form $y = a, z = x + c$, where c is any constant
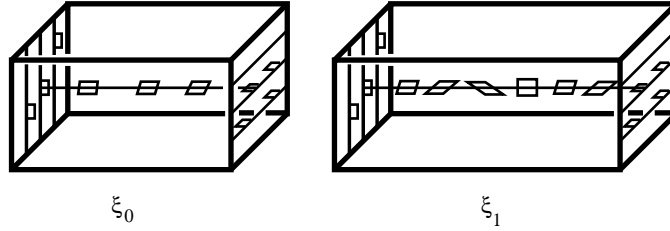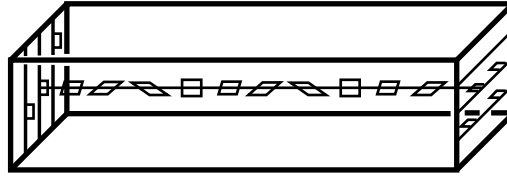
Figure 9: Contact structures ξ_0 and ξ_1.

Figure 10: Contact structure ξ_2.

and $-1 < a < 1$. In addition, there are two circle leaves obtained by performing the identifications at $y = \pm 1$.

Now that we have our contact structures described explicitly, we will choose a transverse 1-foliation Ω_j for each ξ_j to begin our construction of branched surfaces. The 1-foliations are all constructed in the same way. First notice that the induced foliation of ξ_j on the section T_y^2, $-1 < y < 1$, is a foliation by constant slope curves (lines in the model), and the slope rotates from ∞ to 0 through an angle of $\pi/2 + j\pi$. We define a vector field P_j on our cube model tangent to planes $y = $ const and an orthogonal unit vector to the foliation induced on these planes by ξ_j. The vector field induces a vector field on S^3 which we also denote P_j. Let $H = (0, f(y), 0)$ be a horizontal vector field, where f is zero at $y = \pm 1$ and positive elsewhere, chosen to give a smooth vector field on the quotient space. We now let Ω_j be the 1-foliation obtained by integrating $X_j = P_j + H$, see Figure 11. Notice that the figure applies only for ξ_0 and ξ_2, not ξ_1, because P is opposite for ξ_1 at the right end of the model.
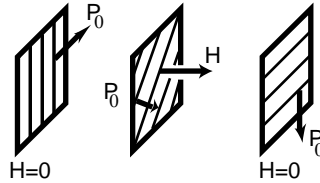

Figure 11: The vector fields P_0 and H.

The reader can verify that Ω_j has two circle leaves, the circles obtained by performing identifications at $y = \pm 1$, and all other (line) leaves limit on both of these closed curves. Regarding the leaves as orbits of the flow determined

33

by the vector field, as $t \to \infty$ the line leaves spiral towards the right circle, as $t \to -\infty$ they spiral to the left circle leaf.

We are now ready to begin constructing branched surfaces. In all cases, we use two discs to split the contact structure. We perform a small σ-deformation of ξ_j to make these integral surfaces. The discs are shown in Figure 12. (The dotted Legendrian joining the two discs in the figure should be ignored for now.) Cutting on the splitting discs yields a manifold V_j with inward cusps. The foliation Ω_j foliates V_j by intervals, giving V_j the structure $V(\hat{B}_j)$ of a fibered neighborhood of a branched surface $\hat{B}_j$.
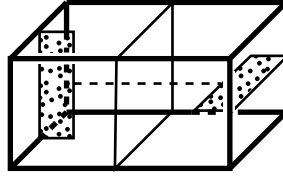

Figure 12: Splitting Discs.

After cutting the leaves of Ω_j on the discs of splitting we get three kinds of segments. Segments of the first kind have both ends on the left cutting disc; these segments are contained in the left region shown in Figure 13. (The figure is valid for ξ_0 and ξ_2, not ξ_1.) Segments of the second kind have both ends on the right cutting disc; these segments are contained in the right region shown in Figure 13. Segments of the third kind have one end on the left cutting disc and one end on the right cutting disc. The segments of the third kind all cross the middle torus $\{y = 0\}$ and are contained in the middle region shown in Figure 13. For ξ_1 the right end of Figure 13 must be modified.
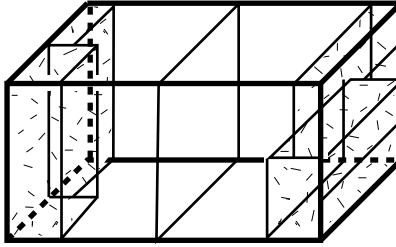

Figure 13: Regions corresponding to sectors of branched surface for ξ_0 and ξ_2.

For each type of segment in the cut Ω_j, we obtain an embedded subsurface of our branched surface $\hat{B}_j$. For segments of the first (or second) kind, we obtain a disc sector in the branched surface, because the segments can be parametrized by their endpoints in the corresponding splitting disc. Every segment of the third kind intersects the torus $y = 0$ exactly once, so these segments yield a torus. To understand how the discs are attached to the torus to form a branched surface, we examine the leaves of Ω_j passing through the boundaries of the regions in Figure 13. Due to the invariance of Ω_j under the action of

34

the group T^2, these are circles, so the discs are attached to the torus on simple closed curves in the torus. To understand the sense of branching in the cases $j = 0, 2$, note that as we move a segment corresponding to a point in the torus up in the model, the corresponding leaf of Ω_j suddenly intersects the splitting disc on the right an additional time. This means the sense of branching where the right disc sector is attached to the central torus is as shown in Figure 14. We denote by $B = \hat{B}_0 = \hat{B}_2$ the branched surface shown in the figure, which is nothing but a double Reeb branched surface in S^3. There is another branched surface $\hat{B}_1 = B_1$ which carries ξ_1, and which has the opposite sense of branching on the right side. The double Reeb branched surface B_1 is shown on the left side of Figure 20.
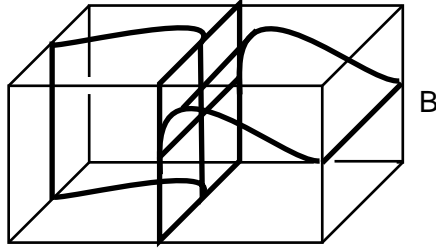

Figure 14: Branched surface carrying ξ_0 and ξ_2.

We are not done because we now have a branched surface B carrying both ξ_0 and ξ_2, one of which is tight and the other overtwisted. In general, our goal is to distinguish tight and overtwisted contact structures using branched surfaces, so this branched surface is not satisfactory. By splitting further, we will obtain two branched surfaces B_0 and B_2, each a splitting of B, such that B_0 carries ξ_0 and B_2 carries ξ_2. The splitting is obtained by joining the two splitting discs with a thin bridge. This bridge can be produced by a σ-deformation performed inside a properly chosen $\mathcal{B}$-chart around the legendrian segment shown in Figure 12 as a dotted segment. A σ-deformation that one can use is, roughly speaking, the opposite of the controlled reduction which eliminates a bridge, as described in Section 2. The resulting splitting discs are shown in Figure 15.
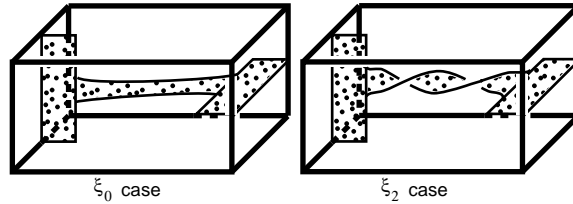

Figure 15: Splitting discs joined to get branched surfaces for ξ_0 and ξ_2.

The effect of extending the splitting surface is to introduce a tunnel joining two parts of the branch locus of B. In the two cases there are two different ways of tunnelling, one to yield B_0 and the other to yield B_2, see figure 16.

For ξ_2, the figure shows the arc along which one tunnels to obtain B_2, joining the two complementary components of B. Another picture of B_2 is shown in Figure 20, together with a picture of B_1, the branched surface carrying ξ_1.
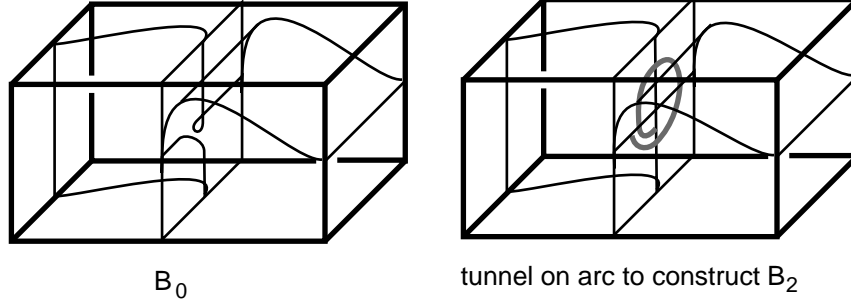

B$_0$ tunnel on arc to construct B$_2$

Figure 16: Branched surfaces for ξ_0 and ξ_2.

We will show in Section 10 that B_1 and B_2 each carry a unique contact structure, which is overtwisted.

8   Pinching and the proof of Theorem 1.10

In this section we deal with the operation of "pinching" for pure contaminations, which is something like the reverse of splitting. We show that using this operation one can recover a contact structure carried by a branched surface B from any pure contamination carried by B representing this structure (Theorem 1.10).

Roughly speaking, a pinching of a pure contamination ξ is an operation which is the composition of the following operations:

(1) identification (or glueing) in a manifold V supporting ξ some diffeomorphic parts of two adjacent components of the horizontal boundary $\partial_h V$, thus getting a new manifold with inward cusps V';

(2) pushing forward ξ to V', which gives the plane field ξ', for which the glueing locus is the maximal integral surface, but which is not necessarily smooth at this locus;

(3) eliminating the integral surface from and simultaneously smoothing ξ'.

These steps require more careful description which we give below.

Let H_1, H_2 be two adjacent components of the horizontal boundary $\partial_h V$. Suppose that for $i = 1, 2$, $F_i \subset H_i$ is a surface (with piecewise smooth boundary) satisfying the following conditions, see Figure 17:

(1) $\partial F_i \subset \mathrm{int}(H_i) \cup (H_1 \cap H_2)$ and the closure of $H_i - F_i$ is a surface (possibly empty) with smooth boundary;

(2) there is a diffeomorphism $h : F_1 \to F_2$ which coincides with each of the germs g_c close to c, for any $c \subset H_1 \cap H_2$ (so that in particular, $F_1 \cap c = F_2 \cap c$).

A diffeomorphism h as above will be called a *glueing diffeomorphism* for V. Strictly speaking, we will always assume that a glueing diffeomorphism h is equipped with a germ of its extension to a neighborhood of F_1 in H_1, and that this extending germ coincides with each of the germs g_c on the overlapping parts.
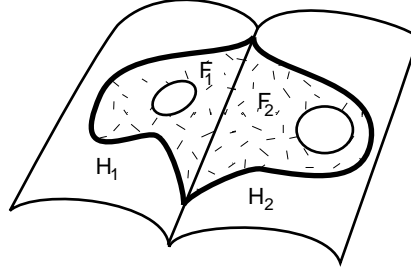

Figure 17: Pinching: identify F_1 and F_2 using glueing diffeomorphism h.

The quotient space V/h with respect to a glueing diffeomorphism h can be equipped with a structure of a smooth manifold with inward cusps, in such a way that the quotient map $V \to V/h$ is smooth. Observe however that there is some freedom of choice for a smooth structure at points of the glueing locus (we will face this problem later in Lemma 8.6). Pushing forward (through the quotient map $V \to V/h$) a pure contamination ξ supported on V, we get a plane field ξ' on V/h, smooth everywhere except possibly at the glueing locus. The behaviour of ξ' at the glueing locus motivates the following.

Definition 8.1. A *semi-smooth σ-contamination* is a plane field ξ on a manifold with inward cusps V which satisfies all the properties of a σ-contamination, except that at the maximal integral surface F it may be semi-smooth (and not just smooth) in the following sense: for a sufficiently small neighborhood U of a point of F, ξ is smooth when restricted to any of the two (closed) sides of U with respect to F. In other words, ξ is semi-smooth if the induced plane field in the manifold V' obtained from V by cutting along the maximal integral surface of ξ is a smooth pure contamination.

Clearly, the plane field ξ' obtained from a pure contamination ξ by a glueing h is a semi-smooth σ-contamination on V/h.

A *semi-smooth σ-deformation* is a 1-parameter family ξ_t of semi-smooth σ-contaminations, smoothly depending on t, which satisfies all other properties of a σ-deformation. It turns out that semi-smooth σ-contaminations and semi-smooth σ-deformations share basic properties with their smooth counterparts. We formulate them in the two results below. Proofs of these results, in principle similar to the smooth case, require a complicated technique of moving charts (extending the technique of special charts as presented in Sections 2 and 3). We will present this technique and sketch the proofs in the next section.

Lemma 8.2. *For any semi-smooth σ-contamination ξ there exists a semi-smooth σ-deformation which turns ξ into a pure contamination. Moreover, for any semi-smooth σ-deformation ξ_t there exists a 1-parameter family $\xi_{t,s}$: $s \in [0,1]$ of semi-smooth σ-deformations (so that $\xi_{t,s}$ depends smoothly on both t and s) which turns $\xi_t = \xi_{t,0}$ parametrically into pure contaminations $\xi_{t,1}$, for all t.*

Proposition 8.3. *Any two contact structures in a closed 3-manifold M which are connected by a semi-smooth σ-deformation, are isotopic.*

We are now ready to give a precise description of an operation of pinching.

Definition 8.4. *Pinching* is an operation which converts a pure contamination ξ_1 into another pure contamination ξ_2 by means of the following two steps:

(1) modifying ξ_1 into a semi-smooth σ-contamination ξ_1' by means of a glueing of the manifold with inward cusps supporting ξ_1 and by pushing forward ξ_1;

(2) modifying ξ_1' to a smooth pure contamination ξ_2 by a semi-smooth σ-deformation (as in Lemma 8.2).

We will use the term "glueing" not only for manifolds V with inward cusps, but also for contaminations ξ with support V.

Observe that the inverse operation of any splitting is an example of pinching, but the inverse of a pinching is not easily seen to be a splitting (this is because each σ-deformation is a semi-smooth σ-deformation, but not vice versa). Nevertheless, we have the following.

Theorem 8.5. *Two contact structures on a closed 3-manifold M which are connected by a sequence of splittings and pinchings are isotopic.*

Notice that Theorem 1.10 of the introduction is a special case of the above theorem. The proof, which is given at the end of the section, requires many preparatory results. We start with a result explaining how the choice of the smooth structure in the supporting manifold influences the semi-smooth σ-contamination obtained by glueing.

Lemma 8.6. *Let ξ be a pure contamination on a manifold with inward cusps V and let h be a glueing diffeomorphism for V. Let $(V/h)_1$ and $(V/h)_2$ be two smooth structures of a manifold with inward cusps for the quotient V/h. Denote by ξ_1 and ξ_2 the semi-smooth σ-contaminations induced by ξ in $(V/h)_1$ and $(V/h)_2$ respectively. Then there is a semi-smooth σ-deformation which turns ξ_1 into ξ_2.*

Proof. Denote by $i : (V/h)_1 \rightarrow (V/h)_2$ the identity map on V/h and notice that it is a semi-smooth isomorphism of ξ_1 and ξ_2, i.e. $i^*(\xi_1) = \xi_2$. On the other hand, $(V/h)_1$ and $(V/h)_2$ are diffeomorphic (as smooth manifolds with inward cusps) through a diffeomorphism $f : (V/h)_1 \rightarrow (V/h)_2$ which is the identity on the glueing locus and which approximates i. Then there exists an

isotopy $f_t : (V/h)_1 \to (V/h)_2$, $t \in [0,1]$, of semi-smooth diffeomorphisms f_t (where the semi-smoothness reflects the bad behaviour at the glueing locus of h), with $f_0 = f$ and $f_1 = i$. Put $\eta_t = f_t^* \xi_1$ and note that $\eta_t : t \in [0,1]$ is a semi-smooth σ-deformation connecting $\eta_0 = f^* \xi_1$ and $\eta_1 = i^* \xi_1 = \xi_2$. Since $f^* \xi_1$ is isomorphic to ξ_1, the lemma follows. $\qquad\square$

The next lemma deals with glueings (of pure contaminations with product complementary pieces) along the entire horizontal boundary, which yield semi-smooth σ-confoliations.

Lemma 8.7. *Let ξ be a pure contamination in a closed 3-manifold M, and suppose that ξ has product complementary pieces. Let h_1, h_2 be two glueing diffeomorphisms that collapse all components of the complement, so that the quotients V/h_1 and V/h_2 are closed topological manifolds homeomorphic to M. Let $(V/h_1)_0$ and $(V/h_2)_0$ be smooth structures in the corresponding quotients, and denote by ξ_1, ξ_2 the semi-smooth σ-confoliations induced from ξ in $(V/h_1)_0$ and $(V/h_2)_0$ respectively. Then ξ_1 and ξ_2 are connected by a semi-smooth σ-deformation.*

Proof. Since both glueing diffeomorphisms h_1, h_2 collapse all product complementary pieces, there is an isotopy $h_t : t \in [1,2]$ between them, through glueing diffeomorphisms h_t. It is possible to choose smooth structures $(V/h_t)_0$ in the quotients V/h_t, for $t \in [1,2]$, so that the corresponding family ξ_t' of semi-smooth σ-confoliations induced by ξ represents a semi-smooth σ-deformation. Since, by Lemma 8.6, the semi-smooth σ-confoliations ξ_i and ξ_i' are connected by a semi-smooth σ-deformation, for $i = 1, 2$, the proof is finished. $\qquad\square$

The next lemma directly prepares the ground for proving Theorem 8.5.

Lemma 8.8. *Let ξ_2 be a pure contamination with product complementary pieces in a closed 3-manifold M, and suppose that it is obtained by pinching from another such pure contamination ξ_1. Let ξ_1' and ξ_2' be any two semi-smooth σ-confoliations in M obtained by glueing from ξ_1 and ξ_2 respectively. Then ξ_1' and ξ_2' are connected by a semi-smooth σ-deformation.*

Proof. Denote by η the semi-smooth σ-contamination obtained after the glueing part of the pinching from ξ_1 to ξ_2 (before eliminating the maximal integral surface and smoothing). Then η and ξ_2 are connected by the semi-smooth σ-deformation occurring in the remaining part of this pinching. We denote this semi-smooth σ-deformation by $\eta_t : t \in [0,1]$, where $\eta_0 = \eta$ and $\eta_1 = \xi_2$. We apply an operation of glueing parametrically to η_t, thus getting η_t', a family of semi-smooth σ-confoliations being a semi-smooth σ-deformation. By Lemma 8.7, there is a semi-smooth σ-deformation between η_1' and ξ_2'. On the other hand, since η_0' is obtained from ξ_1 by the two succesive glueings $\xi_1 \to \eta = \eta_0$ and $\eta_0 \to \eta_0'$, we may view it as obtained from ξ_1 by a single glueing. Applying Lemma 8.7 again, we get that there is a semi-smooth σ-deformation between ξ_1' and η_0'. Composing the above semi-smooth σ-deformations we connect ξ_1' with ξ_2', and the lemma follows. $\qquad\square$

Proof. Theorem 8.5. Suppose that ξ_0, ξ_n are contact structures in a closed 3-manifold M connected by a sequence $\xi_0, \xi_1, \ldots, \xi_{n-1}, \xi_n$ of pure contaminations, so that ξ_{i+1} is obtained from ξ_i either by splitting or by pinching, for $i = 0, 1, \ldots, n-1$. For each $0 \leq i \leq n$, let ξ'_i be a semi-smooth σ-confoliation in M obtained from ξ_i by glueing (we also assume that $\xi'_0 = \xi_0$ and $\xi'_n = \xi_n$). If ξ_{i+1} is obtained from ξ_i by pinching, Lemma 8.8 implies that there is a semi-smooth σ-deformation between ξ'_i and ξ'_{i+1}. The same is true if ξ_{i+1} is obtained from ξ_i by splitting, since then we may view ξ_i as obtained from ξ_{i+1} by pinching. Thus ξ_0 and ξ_n are connected by a semi-smooth σ-deformation. Applying Proposition 8.3 we conclude that the contact structures ξ_0 and ξ_n are isotopic. $\square$

9 Moving charts and properties of semi-smooth σ-deformations

In this section we sketch the proofs of Lemma 8.2 and Proposition 8.3. These proofs go basically along the same lines as proofs of the related results in Sections 2 and 3, where we use special charts corresponding to bridges (or 1-handles) and discs (0-handles) of some handle decomposition of the maximal integral surface. The main difference is that here we not only need to eliminate the maximal integral surface, but also we need to get smoothness from the semi-smoothness in the process. This forces us to work with "moving" charts, and not just with the fixed special charts as in Sections 2 and 3.

To avoid inessential technical complications we will restrict the exposition to the case of semi-smooth σ-confoliations, and skip the details that have to be added in the general case. We start by introducing semi-smooth analogues of special charts.

Let B be a bridge in the maximal integral surface F of a semi-smooth σ-confoliation ξ in a 3-manifold M and let X be a vector field in M transverse to ξ. Recall that $\mathcal{B} = \{(x, y, z) : |x| < 1, |y| < 1, |z| < 1\}$. It is not difficult to get a smooth chart $\varphi : \mathcal{B} \to M$ satisfying the following conditions:

(1) images by φ of the lines $\{x = \text{const}, y = \text{const}\}$ in $\mathcal{B}$ are contained in the integral curves of the vector field X;

(2) $\varphi^{-1}(B) = \{(x, y, z) \in \mathcal{B} : z = 0, |x| \leq 1 - \epsilon, \phi_1(x) \leq y \leq \phi_2(x)\}$ and $\varphi^{-1}(F) = \varphi^{-1}(B) \cup \{(x, y, z) \in \mathcal{B} : z = 0, |x| \geq 1 - \epsilon\}$.

We will call a chart φ as above an *X-chart for B*.

Definition 9.1. With notation as above, a *semi-smooth $\mathcal{B}$-chart for ξ consistent with an X-chart φ for B* is an injective continuous map $\psi : \mathcal{B} \to M$ of the form $\psi(x, y, z) = \varphi(x, y, g(x, y, z))$, with $g(x, y, 0) = 0$, such that

(1) denoting $\mathcal{B}^+ = \{(x, y, z) \in \mathcal{B} : z \geq 0\}$ and $\mathcal{B}^- = \{(x, y, z) \in \mathcal{B} : z \leq 0\}$, the restrictions $\psi^{\pm} = \psi|_{\mathcal{B}^{\pm}}$ of ψ to the sets $\mathcal{B}^{\pm}$ respectively, are smooth embeddings;

(2) the images by ψ of the lines $\{x = \text{const}, z = \text{const}\}$ in $\mathcal{B}$ are tangent to ξ.

Notice that a semi-smooth $\mathcal{B}$-chart as above satisfies all properties of an ordinary $\mathcal{B}$-chart and all properties B of Section 2, except smoothness at the set $\{(x, y, z) \in \mathcal{B} : z = 0\}$. Observe also that a real function $g : \mathcal{B} \to (-1, 1)$ occurring in the definition clearly has to satisfy the following properties:

(1) g depends smoothly on variables x and y;

(2) for all $x, y \in (-1, 1)$, the one variable function $g(x, y, \cdot) : (-1, 1) \to (-1, 1)$ is a continuous injective function whose restrictions to intervals $\{-1 < z \leq 0\}$ and $\{0 \leq z < 1\}$ are smooth embeddings.

The existence of a semi-smooth $\mathcal{B}$-chart for any X-chart φ for a bridge B follows from a construction that requires solving a parametric system of ordinary differencial equations (whose solutions correspond to the curves tangent to ξ mentioned in condition (2) in Definition 9.1). We omit the details.

Semi-smooth $\mathcal{B}$-charts make it possible to express semi-smooth σ-confoliations in terms of slope functions. More precisely, let ξ be a semi-smooth σ-confoliation with maximal integral surface F, and let B be a bridge in F. Let ψ be a semi-smooth $\mathcal{B}$-chart for ξ, consistent with an X-chart φ for B. As in the case of ordinary $\mathcal{B}$-charts, ξ can be expressed in both half-charts $\psi^{\pm}$ (i.e. restrictions of ψ to the subsets $\mathcal{B}^{\pm}$) as the kernel of a (unique) 1-form $dz - f^{\pm}(x, y, z)dx$. We will call the pair $f^{\pm} : \mathcal{B}^{\pm} \to \mathbb{R}$ of real functions as above the *slope function* for ξ in the chart ψ. Note that, because of semi-smoothness of ψ, in general ξ cannot be expressed in terms of a single smooth slope function. Notice also that any slope function $(f^{\pm})$ for a semi-smooth $\mathcal{B}$-chart ψ clearly satisfies the following conditions:

(1) $f^{\pm}(x, y, z) = 0$ for $(x, y, z) \in (\psi^{\pm})^{-1}(F)$;
(2) $\partial f^{\pm}/\partial y(x, y, z) > 0$ for $(x, y, z) \notin (\psi^{\pm})^{-1}(F)$.

In a similar way, one also defines semi-smooth $\mathcal{C}$-charts for disc components of F, and slope functions for these charts.

Our goal is to construct a semi-smooth σ-deformation which eliminates a bridge from the integral surface F of a semi-smooth σ-confoliation ξ. As before, we will do this in terms of a deformation of a slope function. Deformation of a slope function in a $\mathcal{B}$-chart corresponds to rotating a plane field around the integral curves $\{x = \text{const}, z = \text{const}\}$. Since in our semi-smooth case the foliation (of a part of M corresponding to the chart) by these curves is not smooth at F, we cannot (in general) get a smooth plane field by modifying the slope function only. Thus, we have to allow the chart to move, in order to modify the family $\{x = \text{const}, z = \text{const}\}$ of curves into a smooth foliation near the bridge. Simultaneously, modifying the slope function in the moving chart, we will eliminate the bridge from the integral surface. To carry out this plan we start with the following.

Definition 9.2. Let $\psi(x, y, z) = \varphi(x, y, g(x, y, z))$ be a semi-smooth $\mathcal{B}$-chart for ξ consistent with an X-chart φ. A *smoothing X-deformation* for ψ is a 1-

parameter family $\psi_t : t \in [0, 1]$, depending smoothly on t, of maps $\psi_t : \mathcal{B} \to M$ of the form $\psi_t(x, y, z) = \varphi(x, y, g_t(x, y, z))$, such that:

(1) $\psi_0 = \psi$ and for each $t \in [0, 1]$ ψ_t is a continuous embedding;

(2) for each $t \in [0, 1]$ and for all $x, y \in (-1, 1)$ we have $\psi_t(x, y, 0) = \psi(x, y, 0)$, and the restrictions $\psi_t^{\pm}$ of ψ_t to the subsets $\mathcal{B}^{\pm}$ respectively, are smooth embeddings;

(3) for each $t \geq 1/2$ the map ψ_t is smooth at points $(x, y, z) \in \mathcal{B}$ with $z = 0$ and $|x| < 1 - \epsilon + \rho$, for some $0 < \rho < \epsilon$;

(4) there is $\delta > 0$ such that for all $t \in [0, 1]$ the images by ψ_t of the curves $\{x = \text{const}, z = \text{const}\}$ in $\mathcal{B}$ are tangent to ξ at all points $(x, y, z) \in \mathcal{B}$ with $|x| \geq 1 - \delta$ or $|y| \geq 1 - \delta$ or $|z| \geq 1 - \delta$.

The smoothing X-deformation will be our tool for eliminating a bridge from the maximal integral surface of a semi-smooth σ-confoliation. Thus, in condition (3) above we require smoothness for ψ_t with $t \geq 1/2$ in order to get smoothness near the eliminated bridge in modified semi-smooth σ-confoliation. The role of the small constant ρ in this condition is related to the fact that, according to the definition, a semi-smooth σ-confoliation is smooth in some thin neighborhood N of ∂F in the maximal integral surface F. The set $\{(x, y, z) \in \mathcal{B} : z = 0, 1 - \epsilon < |x| < 1 - \epsilon + \rho\}$ will correspond to such a neighborhood N after eliminating the bridge B from the maximal integral surface.

Condition (4) means roughly that the foliation by the images of the curves $\{x = \text{const}, z = \text{const}\}$ does not change near the frontier of the chart neighborhoods in M, when the parameter t runs through $[0, 1]$. This will make it possible to modify ξ inside the chart neighborhoods only. The constant δ occurring in this condition has to be so small that the preimages $\psi_t^{-1}(B) = \psi^{-1}(B)$ of the bridge B are contained in the set $(x, y, z) \in \mathcal{B} : |y| \leq 1 - \delta$; it also has to be smaller than ϵ.

Observe that a smoothing X-deformation ψ_t for ψ is fully determined by a 1-parameter family $g_t : t \in [0, 1]$ of functions $g_t : \mathcal{B} \to (-1, 1)$, smoothly depending on t, satisfying the following conditions:

(1) $g_0 = g$ and for all $t \in [0, 1]$ and all $x, y \in (-1, 1)$ the maps $g_t(x, y, \cdot)$ are continuous and injective;

(2) for all $t \in [0, 1]$ and all $x, y \in (-1, 1)$ we have $g_t(x, y, 0) = g(x, y, 0)$ and $g_t(x, y, \cdot)$ is a smooth embedding when restricted to intervals $\{-1 < z \leq 0\}$ or $\{0 \leq z < 1\}$;

(3) for each $t \geq 1/2$ the function g_t is smooth at $\{(x, y, z) \in \mathcal{B} : z = 0, |x| < 1 - \epsilon + \rho\}$, for some $0 < \rho < \epsilon$;

(4) there is $\delta > 0$ such that for all $t \in [0, 1]$ we have that locally $g_t(x, y, z) = g(x, y, u_t(x, z))$ for some functions u_t and for $|x| \geq 1 - \delta$ or $|y| \geq 1 - \delta$ or $|z| \geq 1 - \delta$.

Note that these conditions correspond to the conditions (1)-(4) in Definition 9.2.

To eliminate a bridge by means of a semi-smooth σ-deformation we need to enrich a smoothing X-deformation corresponding to this bridge with a consistent deformation of the slope function $f^{\pm}$. To do this, we need further restrictions on our smoothing X-deformation. The next definition and lemma explain the details of this idea.

Definition 9.3. With notation as in Definition 9.2 for all related objects, let ψ_t be a smoothing X-deformation. Due to property (4) in this definition, the semi-smooth σ-confoliation ξ can be expressed in charts ψ_t, near their frontiers, by means of slope functions. More precisely, put $\mathcal{B}_{\delta} := \{(x, y, z) \in \mathcal{B} : |x| \geq 1 - \delta, |y| \geq 1 - \delta, |z| \geq 1 - \delta\}$, $\mathcal{B}_{\delta}^+ := \{(x, y, z) \in \mathcal{B}_{\delta} : z \geq 0\}$ and $\mathcal{B}_{\delta}^- := \{(x, y, z) \in \mathcal{B}_{\delta} : z \leq 0\}$. Then there is the unique family $h_t^{\pm} : \mathcal{B}_{\delta}^{\pm} \to \mathbb{R}$ of functions such that ξ is equal to the kernel $\ker[dz - h_t^{\pm}(x, y, z)dx]$ in the chart ψ_t.

We say that ψ_t is *slope-extendable*, if for each $t \in [0, 1]$ and for each (x, z) not in $\{z = 0, |x| \leq 1 - \epsilon\}$ we have $h_t^{\pm}(x, -1 + \delta, z) < h_t^{\pm}(x, 1 - \delta, z)$.

Observe that the inequality as above is automatically satisfied by the slope function $f^{\pm}$ for ξ, due to monotonicity with respect to the variable y. Thus, this inequality is a necessary condition for partial slope functions $h_t^{\pm}$ to be extendable to appropriate slope functions (for the constructed deformation) in the whole $\mathcal{B}$.

As we shall see below, slope-extendability makes it possible to extend functions $h_t^{\pm}$ to a deformation of slope functions associated to ψ_t, which describes a semi-smooth σ-deformation eliminating a bridge from the maximal integral surface. It is thus an important question whether slope-extendable smoothing X-deformations exist. The answer is given in the following.

Lemma 9.4. *Suppose we are given a semi-smooth σ-confoliation ξ, a bridge B in the maximal integral surface for ξ, a vector field X transverse to ξ, and an X-chart φ for B. Then we can find a semi-smooth $\mathcal{B}$-chart ψ for B, consistent with φ, for which there exists a slope-extendable smoothing X-deformation ψ_t.*

We delay the proof of above lemma, and first show how to apply it to eliminate a bridge from a semi-smooth σ-confoliation ξ.

Proposition 9.5. *Let ξ be a semi-smooth σ-confoliation with the maximal integral surface F, and let B be a bridge in F. Then there exists a semi-smooth σ-deformation $\xi_t : t \in [0, 1]$, with $\xi_0 = \xi$, which eliminates B from F, i.e. for which the maximal integral surface of ξ_1 is equal to the closure of $F - B$ and there is only one discontinuity point for the maximal integral surface of ξ_t (where the surfaces changes from F to the closure of $F - B$).*

Proof. Choose a vector field X transverse to ξ and an X-chart φ for B. By Lemma 9.4, there exists a semi-smooth $\mathcal{B}$-chart ψ for ξ consistent with φ, which

43

admits a slope-extendable smoothing X-deformation ψ_t. Let $h_t^\pm : \mathcal{B}_\delta^\pm \to \mathbb{R}$ be the partial slope functions for ψ_t. By the slope-extendability of ψ_t, we can extend $h_t^\pm$ to a family $f_t^\pm : \mathcal{B}^\pm \to \mathbb{R}$ of smooth (pairs of) functions, smoothly depending on $t \in [0,1]$, satisfying the following conditions:

(1) for each $t \in [0, 1/2]$ we have $\partial f_t^\pm/\partial y(x,y,z) > 0$ for $(x,y,z) \in \psi_t^{-1}(M - F)$ and $\partial f_t^\pm/\partial y(x,y,z) = 0$ for $(x,y,z) \in \psi_t^{-1}(F)$;

(2) for each $t \in (1/2, 1]$ we have $\partial f_t^\pm/\partial y(x,y,z) > 0$ for $(x,y,z) \in \psi_t^{-1}(M - cl(F - B))$ and $\partial f_t^\pm/\partial y(x,y,z) = 0$ for $(x,y,z) \in \psi_t^{-1}(cl(F - B))$;

(3) for each $t \in [0,1]$ the 1-forms $dz - f_t^\pm(x,y,z)dx$ in $\mathcal{B}^\pm$, after pushing them forward through ψ_t, determine a well-defined 1-form ω_t in $\psi_t(\mathcal{B})$ (i.e. the two pushed-forward 1-forms on the sets $\psi_t(\mathcal{B}^\pm)$ respectively coincide on the intersection of these sets, thus giving a sort of semi-smooth 1-form);

(4) for each $t \in [0, 1/2]$ the 1-form ω_t is smooth at all points $\psi_t(x,y,0)$ outside F or in a small neighborhood N of ∂F in F;

(5) for each $t \in (1/2, 1]$ the 1-form ω_t is smooth at all points $\psi_t(x,y,0)$ outside $cl(F - B)$ or in a small neighborhood N of $\partial[cl(F - B)]$ in $cl(F - B)$.

In fact, slope-extendability is necessary (and sufficient) for conditions (1) and (2) only; the other conditions can be fulfilled without referring to this property. In particular, conditions (4) and (5) can be achieved due to the smoothness of the foliations by the images under ψ_t of the curves $\{x = \text{const}, z = \text{const}\}$ in $\mathcal{B}$, at the corresponding points.

Conditions (1)-(5) above imply that putting

$$\xi_t(p) = \begin{cases} \xi(p) & \text{for } p \notin \psi_t(\mathcal{B}) \\ \ker \omega_t & \text{for } p \in \psi_t(\mathcal{B}) \end{cases}$$

we get a semismooth σ-deformation ξ_t that eliminates the bridge B as required, thus finishing the proof. $\square$

Proof. (Lemma 9.4.) Choose a semi-smooth $\mathcal{B}$-chart ψ consistent with φ so that it satisfies the following additional condition: images by ψ of the lines $\{y = 0, z = \text{const}\}$ in $\mathcal{B}$ are everywhere tangent to ξ. It is not difficult to convince oneself that this can always be done. Note that in such a $\mathcal{B}$-chart ψ the slope function $f^\pm$ for ξ satisfies the condition: $f^\pm(x, 0, z) = 0$ for all $x, z \in (-1, 1)$. Likewise, due to monotonicity of $f^\pm$ with respect to y outside the integral surface F, we have $f^\pm(x, 1 - \delta, z) > 0$ and $f^\pm(x, -1 + \delta, z) < 0$ for any (x, z) outside $\{z = 0, |x| \geq 1 - \epsilon\}$. We will show the existence of a slope-extendable smoothing X-deformation ψ_t for any semi-smooth $\mathcal{B}$-chart ψ as above.

We will concentrate on constructing the appropriate final map ψ_1 for a slope-extendable smoothing X-deformation $\psi_t : t \in [0,1]$. We will comment on constructing the whole family $\psi_t : t \in [0,1]$ only after constructing ψ_1.

Choose rectangles $R_+ := \{y = 1-\delta\}$, $R_- := \{y = -1+\delta\}$ and $R_0 := \{y = 0\}$ in $\mathcal{B}$, and view them as subsets of M expressed in the chart ψ. Consider

characteristic foliations $\zeta_+, \zeta_-, \zeta_0$ of ξ at R_+, R_-, R_0 respectively. Then, due to additional condition on ψ as above, leaves of ζ_0 all have the form $\{z = \text{const}\}$. Moreover, again due to the monotonicity with respect to y, the slopes of ζ_+ (ζ_-) are positive (negative respectively) everywhere except at $\{z = 0, |x| \geq 1 - \epsilon\}$, where they are zero, see Figure 18.

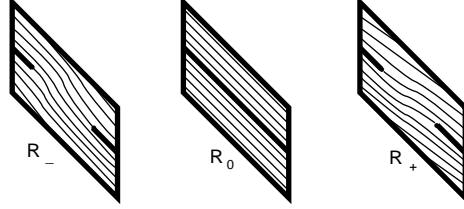
Figure 18: Characteristic foliation on rectangles.

Consider now smooth coordinates x', y', z' in $\psi(\mathcal{B})$ with $x' = x$, $y' = y$ and $z' = z'(x, y, z)$, such that $\{z = 0\} = \{z' = 0\}$ (i.e. $z'(x, y, 0) = 0$). In these coordinates ζ_0 is smooth everywhere except at $\{z' = 0\}$, while ζ_+ and ζ_- are smooth everywhere except at $\{z' = 0, |x'| \geq 1 - \epsilon + \rho\}$. Let χ_+ and χ_- be the 1-foliations in R_+, R_- respectively, with all leaves of form $\{z = \text{const}\}$ in the initial chart ψ. Then the slopes of χ_+ (χ_-) are smaller (respectively bigger) than the slopes of ζ_+ (ζ_- respectively), except at $\{z' = 0, |x'| \geq 1 - \epsilon\}$, where the slopes of both foliations are equal to zero (horizontal). Moreover, the foliations χ_+ and χ_- expressed in (x', z')-coordinates are semi-smooth (smooth when restricted to $\{z' \geq 0\}$ and to $\{z' \leq 0\}$, but in general not smooth at $\{z' = 0\}$).

A crucial step in constructing a chart ψ_1 consists of replacing the foliations χ_+ and χ_- by foliations χ'_+ and χ'_- respectively, satisfying the following conditions:

(1) χ'_+ (χ'_-) coincides with χ_+ (χ_- respectively) at the subset in R_+ (R_- respectively) described as $\{|x| \geq 1 - \delta$ or $|z| \geq 1 - \delta\}$ in the coordinates (x, z) from the chart ψ;

(2) the curve $\{z' = 0\}$ in R_+ (R_-) is still a leaf of χ'_+ (χ'_- respectively);

(3) the foliations $\chi'_{\pm}$ are smooth everywhere except possibly at $\{z' = 0, |x'| \geq 1 - \epsilon + \rho\}$;

(4) χ'_+ (χ'_-) is still transverse to (the restriction to R_+, or R_- respectively, of) the vector field X, and its slopes are still smaller (bigger respectively) than the slopes of ζ_+ (ζ_-), except at $\{z' = 0, |x'| \geq 1 - \epsilon\}$, where these slopes are equal.

Without going into complete detail, we observe that foliations χ'_+, χ'_- as described above exist by C^0-approximation of a continuous foliation by a smooth one (here we view foliations as line fields).

Choose a 2-foliation $\mathcal{F}$ in $\psi(\mathcal{B})$ transverse to the vector field X such that:

(1) leaves of χ'_+ and χ'_- are contained in the leaves of $\mathcal{F}$;

45

(2) images by ψ of the pieces of lines $\{x = \mathrm{const}, z = \mathrm{const}\}$ contained in $\mathcal{B}_\delta$ are contained in the leaves of $\mathcal{F}$;

(3) $\mathcal{F}$ is smooth except at $\{z' = 0, |x'| \geq 1 - \epsilon + \rho\}$.

Such a foliation $\mathcal{F}$ can easily be constructed due to the properties of the previously constructed foliations χ'_+ and χ'_-.

Choose a transverse parameter with respect to which $\mathcal{F}$ is smooth as in (3), and take this parameter as z-coordinate of the chart ψ_1, keeping the x- and y-coordinates as in ψ. Then, due to conditions (1)-(3) above, the chart ψ_1 and its partial slope function $h_1^\pm$ fulfill the relevant requirements from the definition of a slope-extendable smoothing X-deformation. In particular, condition (1) above and condition (4) for χ'_+ and χ'_- imply that $h^\pm(x, 1 - \delta, z) > 0$ and $h^\pm(x, -1 + \delta, z) < 0$ for any (x, z) outside $\{z = 0, |x| \geq 1 - \epsilon\}$, hence the slope-extendability of ψ_1.

A construction as above for ψ_1 can be performed parametrically, starting with constructing a parametric family $(\chi'_\pm)_t$ of modified foliations, smoothly connecting $(\chi'_\pm)_0 = \chi_\pm$ with $(\chi'_\pm)_1 = \chi'_\pm$. After taking care for all the necessary details one gets a slope-extendable smoothing X-deformation as required. We omit further details. $\qquad\square$

We now conclude the proof of Lemma 8.2 and Proposition 8.3. To prove the first one, it is sufficient to note that, similarly to the above, one can define and use moving charts for semi-smooth $\mathcal{C}$-charts in order to construct semi-smooth σ-deformations that eliminate a disc component from the maximal integral surface. Extending these methods slightly, one can also get a parametric elimination of the maximal integral surface from any semi-smooth σ-deformation, again using an argument similar to the one in Section 2.

To prove Proposition 8.3 observe that we can parametrically eliminate the maximal integral surface in such a way, that whenever this surface is initially empty, it remains empty during the elimination process. This clearly leads to an isotopy between any two contact structures connected by a semi-smooth σ-deformation, which finishes the proof.

10 Pseudo-transversal knots

In this section we discuss in detail pseudotransversal knots for branched surfaces embedded in 3-manifolds. We start with the definition.

Definition 10.1. Let $B \hookrightarrow M$ be a branched surface embedded in a 3-manifold M, with product complementary pieces. A *pseudotransversal knot* for B is a null-homologous knot k in M satisfying the following conditions:

(1) k is disjoint with the branched locus of B;

(2) k intersects B transversely;

(3) for each connected component U of the complement $M - B$ the intersection $k \cap U$ is (isotopic rel boundary ∂U to) a *braid* in U.

Condition (3) above needs more detailed explanation. We say that an intersection $k \cap U$ is a *braid* in U, if for some identification of U with the product $\text{int}F \times (0, 1)$ each strand (i.e. component) of $k \cap U$ projects to the factor $(0, 1)$ in the product injectively.

Recall now that for a null-homologous knot k transverse to a plane field ξ in a 3-manifold M one defines the self-linking number $l(k)$ (relative to ξ) as follows. Let Σ be a Seifert surface for k in M. Then we can find a non-zero section σ of $\xi|_\Sigma$. We then use σ to get a parallel copy k' of k by pushing towards σ. The self-linking number $l(k)$ is then the linking number of k and k'. One easily verifies that this number does not depend on the choices of Σ and σ (compare [3] or [5]).

If k is a pseudotransversal knot for B, one can define the self-linking number $l(k)$ (relative to B) by means of a plane field ξ_B in M associated to B as follows. At any point of B define ξ_B to be equal to the tangent plane to B. Extend ξ_B to the components U of $M - B$ so that, identifying U with $\text{int}F \times (0, 1)$, restriction of ξ_B to U coincides with the foliation of U by the surfaces $\text{int}F \times \{z\} : z \in (0, 1)$. This clearly defines a continuous plane field in M (we do not require smoothness here), and this field is well-defined up to homotopy of continuous plane fields. Observe now that, according to the definition of a pseudotransversal knot, a plane field ξ_B as above can be chosen in such a way that k is everywhere transverse to ξ_B. The self-linking number $l(k)$ relative to ξ_B is then uniquely determined (i.e. does not depend on the choice of ξ_B), and we take it as the *self-linking number for the pseudotransversal knot k relative to B.*

We now turn to proving the following crucial lemma which directly implies Theorem 1.13 of the introduction, by referring to the Bennequin inequality, mentioned there.

Lemma 10.2. *Suppose that a branched surface $B \hookrightarrow M$ carries a contact structure ξ, and let k be a pseudotransversal knot for B. Then there exists a knot $\tilde{k}$ which is everywhere transverse to ξ, such that $l(\tilde{k}) = l(k)$ (where the self-linking of $\tilde{k}$ is viewed relative to ξ).*

Proof. Let ξ_1 be a pure contamination supported by $V(B)$ obtained from ξ by a splitting. Then k determines a knot k_1 which intersects $V(B)$ along vertical fibres (which correspond to points of transverse intersection of k with B) and whose intersections $k_1 \cap U$ with connected components of $M - V(B)$ are the same braids as the corresponding intersections of k with $M - B$.

Fix a component U of the complement $M - V(B)$ and suppose its boundary consists of the two copies of a surface F, so that U is homeomorphic to $\text{int}F \times (0, 1)$. Suppose that the braid $k_1 \cap U$ consists of m strings. This braid can be then viewed as an isotopy of the set of m points in $\text{int}(F)$. This isotopy can be extended to an isotopy of F in itself (rel boundary), i.e., a 1-parameter family of diffeomorphisms of $f_t : F \to F$, starting at id_F and equal the identity on ∂F

47

at all times. We denote by f_U the final diffeomorphism of such a 1-parameter family. Consider then an operation of pinching that collapses U, with f_U taken as a glueing diffeomorphism. Observe that after doing such pinchings for all connected components of $M - V(B)$, the pieces of $k_1 \cap V(B)$ glue together to form a knot k_2 transverse to the contact structure ξ_2 obtained by these pinchings.

It follows from the description above that the self-linking $l(k_2)$ relative to ξ_2 is equal to the self-linking $l(k)$ relative to B. One can see this by constructing a homotopy between the plane field ξ_B associated to B and the contact structure ξ_2, followed by a homotopy between k and k_2 through knots transverse to the homotoped plane fields. On the other hand, according to Theorem 8.5, ξ_2 is isomorphic to the initial contact structure ξ, and thus the lemma follows. $\qquad\square$

Our next observation is the following.

Lemma 10.3. *Suppose that $B \hookrightarrow M$ is a branched surface that carries a contact structure ξ, and let ξ_B be a continuous plane field in M associated to B (as above, in the definition of the self-linking relative to B). Then ξ and ξ_B are homotopic as plane fields.*

Proof. Let ξ_s be a pure contamination carried by $V(B)$ obtained from ξ by a splitting. We extend ξ_s to a continuous plane field $\xi_{s,B}$ just as we extended the tangent plane field of B to ξ_B. It is then clear that $\xi_{s,B}$ is homotopic both to ξ (by collapsing the product components of $M - V(B)$, which corresponds to pinching ξ_s to return to ξ) and to ξ_B (by collapsing $V(B)$ along vertical fibres), and thus the lemma follows. $\qquad\square$

We summarize above results in the following.

Proposition 10.4. *Let $B \hookrightarrow M$ be a branched surface with product complementary pieces and suppose that it admits a pseudotransversal knot k which violates the Bennequin inequality. Then B carries at most one (up to isotopy) contact structure which, if it exists, is overtwisted.*

Proof. According to Theorem 1.13, any contact structure carried by B is overtwisted. On the other hand, it follows from Lemma 10.3 that any two contact structures carried by B are homotopic as plane fields in M. By the classification theorem for overtwisted contact structures ([2]), there is only one overtwisted contact structure in each homotopy class of plane fields in M, hence the Proposition. $\qquad\square$

The next result, whose proof uses pseudotransversal knots, strengthens Theorem 1.9 in the case of overtwisted contact structures.

Proposition 10.5. *For any overtwisted contact structure ξ in a 3-manifold M there exists a branched surface $B \hookrightarrow M$ such that ξ is the only contact structure it carries.*

Proof. To construct a branched surface B as required we refer to the construction from the proof of Theorem 1.9 in Section 6, adding appropriate modifications.

Let k be a knot in M transverse to ξ and violating the Bennequin inequality (since ξ is overtwisted, such a knot always exists). Let Ω be a 1-foliation in M transverse to ξ, for which k is a leaf. By applying to Ω the construction from the proof of Theorem 1.9 in Section 6 we get a branched surface B for which k is a pseudotransversal knot, with the self-linking relative to B equal to its self-linking relative to ξ and thus violating the Bennequin inequality. It then follows from Proposition 10.4 that ξ is the only contact structure carried by B, which finishes the proof. $\square$

We will now describe an easily detected property of branched surfaces which implies the existence of a pseudotransversal knot that violates the Bennequin inequality.

Definition 10.6. Given a branched surface $B \hookrightarrow M$, an *embedded disc negatively twisted on one side* in B is a disc D smoothly embedded in B, with ∂D transverse to the branch locus of B, with distinguished side, such that all branching of B at branch curves meeting ∂D on the distinguished side of D have the same sense, a sense we call negative, as shown in Figure 19. Observe that if all the branching as above has the same sense, but different from the one in Figure 19, then D is not negatively twisted on the distinguished side. Observe also that for a disc negatively twisted on one side, the branching on the other side can be arbitrary.
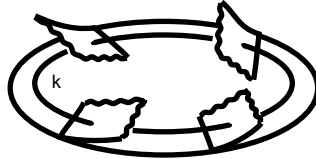

Figure 19: Twisted disc.

Lemma 10.7. *Suppose that a branched surface $B \hookrightarrow M$ with product complementary pieces contains an embedded disc D negatively twisted on one side. Then there exists an unknotted curve k pseudotransversal to B with $lk(k) = 1$. Since this curve violates the Bennequin inequality, the branched surface B carries at most one contact structure which, if it exists, is overtwisted.*

Proof. In view of Proposition 10.4, it is sufficient to prove the first assertion. A curve k as asserted is shown at Figure 19, being a small shift of ∂D towards the distinguished side. It can be checked directly that $lk(k) = 1$. $\square$

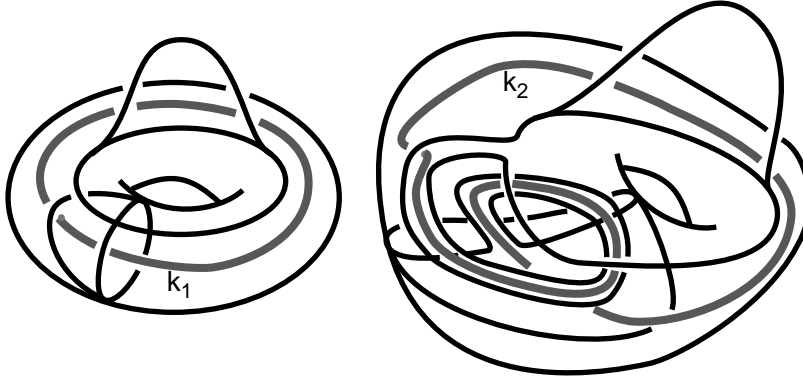As an application of Proposition 10.4 to concrete examples we give the following.

Figure 20: Pseudo-transversals in S^3.

Proposition 10.8. *The branched surfaces B_1, B_2 in S^3 that carry the overtwisted contact structures ξ_1, ξ_2 respectively, as described in Section 7, both carry no other contact structures.*

Proof. Figure 20 shows the pseudotransversal unknotted curves k_1 and k_2 for the branched surfaces B_1 and B_2 respectively, for which we have $lk(k_1) = 1$ and $lk(k_2) = 1$. Since these curves violate the Bennequin inequality, the proposition follows. $\qquad\square$

11 Invariance of tightness for pure contaminations

In the remainder of the paper we shift our attention to contaminations. The goal of this section is to clarify the notion of tightness for pure contaminations and to establish basic classes of deformations of pure contaminations that preserve tightness. This is important for our long-term goal of using tight contaminations to study 3-manifolds.

We start this section by showing that the definition of tightness of a pure contamination as given in the introduction (Definition 1.14) can be simplified slightly.

Given a manifold with inward cusps V embedded in a 3-manifold M, a *compressing disc* for the horizontal boundary $\partial_h V$ is a disc D embedded in M, with $\partial D \subset \partial_h V$ and D transversal to $\partial_h V$ at ∂D, for which there is no disc $H \subset \partial_h V$ with $\partial H = \partial D$. We will say that the horizontal boundary $\partial_h V$ is *incompressible* in M if there is no compressing disc for it. Note that, according to this definition, an overtwisting disc D with $\partial D \subset \partial_h V$ for a pure contamiantion ξ supported by V is a compressing disc for $\partial_h V$. Thus, if ξ is tight then the horizontal boundary $\partial_h V$ of the supporting manifold V is incompressible.

Lemma 11.1. *Suppose ξ is a pure contamination supported by V. Then ξ is tight if and only if there is no overtwisting disc D with $\partial D \hookrightarrow int(V)$ and no compressing disc for $\partial_h V$.*

Proof. We will show that an overtwisting disc D whose boundary ∂D intersects $\partial_h V$, but is not contained in it, can be always replaced by another overtwisting disc D' with $\partial D' \subset int(V)$, which clearly implies the lemma.

For a disc D as above, we can subdivide the boundary ∂D into consecutive segments $I_1, I_2, \ldots, I_{2n}$ such that:

(1) the intersection of ∂D with $\partial_h V$ is contained in the union of segments I_{2k} with even indices only;

(2) for each $1 \leq k \leq n$ the segment I_{2k} intersects $\partial_h V$ in one component only.

We can then push each segment I_{2k} slightly, in a direction transverse to ξ, to get a legendrian segment $I'_{2k} \subset int(V)$. If these modifications are slight enough, we can also perturb the remaining segments I_{2k-1}, which all pass through $int(V)$, into new legendrian segments $I'_{2k-1} \subset int(V)$, so that the union of all segments $I'_j : 1 \leq j \leq 2n$ forms a smooth closed legendrian curve in $int(V)$. It is clear that this curve still bounds an overtwisting disc. □

We now pass to the discussion of deformations (of pure contaminations) that preserve tightness.

Definition 11.2. A *pure deformation away from $\partial_h V$* is a pure deformation ξ_t of (pure) contaminations supported by V such that for some open neighbourhood U of $\partial_h V$ in V the deformation ξ_t is constant on U.

The next lemma is a well-known generalization of Gray's Theorem, and we will refer to it as the relative Gray's Theorem.

Lemma 11.3. *Let $\xi_t : t \in [0,1]$ be a pure deformation away from $\partial_h V$. Then ξ_t is an isotopy and in particular the pure contaminations ξ_0 and ξ_1 are contactomorphic by a diffeomorphism isotopic to identity.*

Proof. An isotopy as in the conclusion of Gray's Theorem can be constructed as a flow with respect to some variable vector field X_t determined by ξ_t, see [1]. This method can be applied in our situation without any changes, giving a variable vector field X_t in V which vanishes near $\partial_h V$. The flow with respect to X_t is then equal to identity in the neighbourhood U of $\partial_h V$ in V, thus giving a 1-parameter family f_t of diffeomorphisms of V. The fact that f_t is an isotopy, as required, follows by arguments described in [1] for the proof of the standard Gray's Theorem. □

The next definition and lemma generalize the above considerations slightly.

Definition 11.4. A *σ-deformation away from ∂V* is a σ-deformation ξ_t of contaminations supported by V such that for some open neighbourhood U of ∂V in V the deformation ξ_t is constant on U.

Observe that if a σ-deformation ξ_t as above starts at a pure contamination, then the maximal integral surface is disjoint from ∂V for all t, i.e. none of its components is adjacent to $\partial_h V$ at the cusp locus. The next lemma shows that tightness is invariant under such deformations.

Lemma 11.5. *Let $\xi_t : t \in [0, 1]$ be a σ-deformation away from $\partial_h V$ connecting pure contaminations ξ_0 and ξ_1. Then ξ_0 and ξ_1 are contactomorphic by a diffeomorphism isotopic to identity. In particular, ξ_0 is tight iff ξ_1 is tight.*

Proof. For a σ-deformation ξ_t away from ∂V it is possible to construct a parametric family $\xi_{t,s}$ of combinations of controlled reductions, which are all away from $\partial_h V$, such that $\xi_{t,0} = \xi_t$ and $\xi_{t,1}$ is pure for all $t \in [0, 1]$. This can be done by using the same methods as in Sections 2 and 3. It follows that the pure contaminations $\xi_0 = \xi_{0,0}$, $\xi_{0,1}$, $\xi_{1,1}$ and $\xi_{1,0} = \xi_1$ are connected by a pure deformation away from ∂V. We finish the proof by applying the relative Gray's Theorem (Lemma 11.3). $\square$

We now turn to deformations that change the supporting manifold V of a pure confoliation. To define a class of such deformations which we will be interested in, recall that a controlled reduction is a kind of σ-deformation, performed by modifying slope functions in a given special chart (a $\mathcal{C}$-chart or a $\mathcal{B}$-chart), with one discontinuity point for the maximal integral surface F, where a disc component or a bridge (1-handle) is eliminated from F (compare Section 2). Recall also that in the context of σ-contaminations, a bridge that is eliminated by a controlled reduction may be partly adjacent to the cusp locus of the supporting manifold V.

Definition 11.6. An *elementary controlled splitting* is a deformation that turns a pure contamination ξ into another pure contamination ξ', and consists of the following two steps. The first step is an operation inverse to a controlled reduction (which in our case either produces a disc component of the maximal integral surface in the interior of V, or produces a bridge adjacent to the cusp locus of V along two disjoint arcs). A second step consists of cutting V along the whole maximal integral surface (produced in the first step), to get a new V' and ξ'. A *controlled splitting* is a finite sequence of elementary controlled splittings. A *controlled pinching* is an operation inverse to a controlled splitting. Finally, a *controlled deformation* is a combination of controlled splittings and pinchings (in fact, it is a combination, in sequence, of elementary controlled splittings and their inverses).

Lemma 11.7. *Controlled deformations of pure contaminations that preserve the property of incompressibility for horizontal boundaries of the supporting manifolds with inward cusps preserve tightness.*

Proof. According to Lemma 11.1, preserving tightness is equivalent to preserving the following two conditions:

(1) nonexistence of an overtwisting disc D with $\partial D \subset int(V)$;

(2) incompressibility of $\partial_h V$.

Under assumptions of this lemma, condition (2) is preserved automatically. It is then sufficient to show that condition (1) is also preserved, and this can be further reduced to the case of a single elementary controlled splitting. We will deal with this case using the methods of Section 4 as follows.

Suppose that $\xi_t : t \in [0,1]$ is the operation inverse to the controlled reduction occurring as a first step in a given special elementary controlled splitting. Then ξ_0 is a pure contamination and ξ_1 has the maximal integral surface F_1 of one of the two forms mentioned Definition 11.6. Both ξ_0 and ξ_1 are supported by the same manifold with inward cusps V. An argument as in Section 4 shows that the existence of an overtwisting disc D for ξ_1 with $\partial D \subset int(V) \setminus F_1$ is equivalent to the existence of such a D for ξ_0 with $\partial D \subset int(V)$.

Let ξ_1' be the pure contamination (supported by V') obtained from ξ_1 by cutting V along F_1, i.e. the one obtained from ξ_0 by our fixed special elementary controlled splitting. From what was said above, it immediately follows that the existence of an overtwisting disc D for ξ_1' with $\partial D \subset int(V')$ is equivalent to the existence of such a D for ξ_0 with $\partial D \subset int(V)$. This finishes the proof. $\square$

We finish this section by remarking that we have not proved the invariance of tightness for the following classes of deformations of pure contaminations:

(1) splittings that are not controlled;

(2) pinchings that are not inverse to appropriate splittings;

(3) splittings which do not preserve the incompressibility property of $\partial_h V$.

More work is required to understand invariance of tightness for pure contaminations.

12 Purification

In many situations, given a (positive) confoliation or contamination ξ defined on M or $V \hookrightarrow M$, which is non-integrable on some open set in its support, it is possible to find a deformation which converts ξ to a contact structure or a pure contamination. Such a deformation is called a *purification*. The idea of purification is one of the important ideas in the monograph of Eliashberg and Thurston, see [4], though they cannot be blamed for our terminology. It would be useful to establish an optimal result indicating when a contamination can be purified. We will need to purify contaminations, but only in very special situations. In fact, the controlled reductions of Section 2 constitute a special class of purifications. In the paper [12], we purify using $\mathcal{C}$-charts ($\mathcal{B}$ charts), under the assumption that the contamination is pure adjacent to the $\mathcal{C}$-chart (the ends of the $\mathcal{B}$-chart), just as in the case of controlled reductions.

In Section 13 we will need to purify in a particular situation, which we now describe. Let $B \hookrightarrow M$ be a branched surface with boundary, properly embedded in a 3-manifold M, i.e. an embedding $(B, \partial B) \hookrightarrow (M, \partial M)$. A

fibered neighbourhood $V(B)$ is associated to B as before, with $V(\partial B)$ the corresponding associated fibered neighbourhood of the boundary train track. Let ξ be a confoliation carried by B, i.e. a confoliation in $V(B)$ that is transverse to the vertical fibers of $V(B)$ and tangent to the horizontal boundary $\partial_h V(B)$. Let $R \subset B$ be a rectangle embedded in B, disjoint from the branch locus of B, with two opposite sides contained in ∂B. Furthermore, let $\{(x,y) : |x| \leq 1, |y| \leq 1\}$ be smooth coordinates for R so that the sides $|x| = \pm 1$ are contained in ∂B. Given ϵ, $0 < \epsilon < 1$, put $R_\epsilon = \{|y| < 1 - \epsilon\} \subset R$ and denote by $V(R) = \pi^{-1}(R)$, $V(R_\epsilon) = \pi^{-1}(R_\epsilon)$ the corresponding restrictions of the fibration $\pi : V(B) \to B$. Finally, suppose that for some $0 < \epsilon < 1$ the confoliation ξ is pure in $V(R) - V(R_\epsilon)$ (except at top and bottom where it is tangent to $\partial_h V(B)$), while it can be arbitrary inside $V(R_\epsilon)$.

Lemma 12.1. *Under the assumptions described above, there exists a confoliation ξ' carried by B that coincides with ξ outside $V(R)$ and is pure in $V(R)$.*

Proof. We construct a confoliation ξ as required in two steps: we first construct a suitable variant of a $\mathcal{B}$-chart; then we modify ξ by changing the slope function inside the $\mathcal{B}$-chart. The kind of $\mathcal{B}$-chart we want to use is a smooth map $\psi : [-1,1]^3 \to V(B)$ such that conditions (1) and (2) of Definition 2.3 are satisfied, as are the following:

(1) the first two coordinates x, y in $[-1,1]^3$ coincide with the coordinates x, y in R, in the sense that $\pi \circ \psi(x,y,z) = (x,y)$ in these coordinates;

(2) the chart image $\psi([-1,1]^3)$ coincides with the preimage $\pi^{-1}(R) = V(R)$;

(3) the images by ψ of the segments $\{x = \text{const}, z = \text{const}\} \subset [-1,1]^3$ are tangent to ξ.
The existence of a chart as above is as obvious as the existence of the ordinary $\mathcal{B}$-charts described in section 2.

Let $f : [-1,1]^3 \to \mathbb{R}$ be the slope function for the contamination ξ in the $\mathcal{B}$-chart coordinates as above. We then have $\partial f/\partial y \geq 0$ and $\partial f/\partial y(x,y,z) > 0$ for $|y| > 1 - \epsilon$ and $|z| < 1$. Thus if $1 > y_1 > 1 - \epsilon$ then $f(x, y_1, z) > f(x, -y_1, z)$ for all $x, z \in (-1, 1)$. It follows that there exists a smooth function $f' : [-1,1]^3 \to \mathbb{R}$ with the following properties:

(1) f' coincides with f in the union of the subsets $\{|y| \geq y_1\}$ and $\{|z| = 1\}$ in $\mathcal{B}$;

(2) $\partial f'/\partial y > 0$ in the subset $\{-1 < z < 1\}$.

Since the function f' coincides with f near the frontier of the chart image in $V(B)$, it can be used as a slope function for a new contamination ξ' which coincides with ξ outside the chart image. By Lemma 2.4, ξ' is contact in the subset $\{-1 < z < 1\}$. □

13 Internal tightness

In this section we deal with a property of pure contaminations weaker than tightness, namely a property we call internal tightness. We find some conditions on a branched surface embedded in 3-manifold that imply the internal tightness of any pure contamination carried by the branched surface. As a corollary we get Theorem 1.15. In the next section we will show that in some cases it is possible to deduce the tightness of a pure contamination from its internal tightness. This will make it possible to describe conditions on a branched surface sufficient to imply tightness of any pure contamination carried by the branched surface.

Definition 13.1. A pure contamination ξ supported by V is *internally tight* if it has no overtwisting disc contained in $\text{int}(V)$. It is *internally universally tight* if there is no overtwisting disc contained in the universal cover $\text{int}(\tilde{V}))$ for the induced (lifted) contamination $\tilde{\xi}$.

Clearly, universal tightness implies tightness since an overtwisting disc (if exists) can be always lifted to the universal cover. In this section we will prove this stronger property of internal universal tightness for some classes of pure contaminations.

To formulate our main result in this section we introduce two concepts necessary for describing the assumptions. A closed (smooth) curve α in the branch locus of a branched surface B is *essential* if its embedding in B is a π_1-injective map. Equivalently, α is essential if any of its lifts to the universal cover $\tilde{B}$ of B is a (non-closed) line.

Before introducing a second concept, the concept of an immersed twisted disc of contact, we give some preliminary definitions.

Given a surface $E \hookrightarrow \partial_h V(B)$ for some branched surface $B \hookrightarrow M$, $\pi : V(B) \to B$ defines a map of E to B. Pulling back the branch locus of B to E, we obtain a *T-pattern* on E. It consists of a 1-complex of smooth curves possibly joined at *T-junctions*, see Figure 21. In addition to this 1-complex, we also consider as part of the T-pattern the data giving the sense of branching on each smooth curve of the T-pattern. This is indicated by an arrow transverse to the curve of the T-pattern as shown in the figure. The transverse arrow describes the sense of branching consistent with our definition, given in the introduction, of inward and outward branching along the boundary of a sector. Recall that if sectors W, X, Y are adjacent along an arc γ of branch locus, and if $W \cup Y$ and $X \cup Y$ are smooth, we say that branching along the arc $\gamma \subset \partial Y$ is *inward* for Y and *outward* for X. We indicate the sense of branching using a transverse arrow pointing into Y.

T-patterns are defined in greater generality. Let B be a branched surface and $V(B)$ its fibred neighbourhood, and suppose that $i : E \to B$ is an immersion of a surface E, usually a disc in this paper. Then i determines an immersion $\hat{i} : E \to int(V(B))$ such that $i = \pi \circ \hat{i}$, unique up to homotopy along the fibres of $V(B)$. If $\hat{i}$ is an embedding in $V(B)$, then cutting $V(B)$ on E

yields a neighborhood $V(B')$ with two components of horizontal boundary E_- and E_+ homeomorphic to E. The T-patterns induced on E_- and E_+ are the *T-patterns* induced on the disc E on each side of E respectively. Even when $\hat{i}$ is not an embedding, we can pull back two T-patterns to E, for example by splitting locally.

Suppose the T-pattern on $\partial_h V(B)$ includes a sequence of segments $\gamma_0, \gamma_1 \ldots \gamma_n$ which can be directed so the final end of γ_i is attached to a point in the interior of $\gamma_{i+1}, \mod n$, forming a cycle that bounds a disc (see Figure 21(b)). We call this disc a *twisted disc of contact in the T-pattern*, which is *positive* (*negative*) if the double points of the branch locus at the ends of the γ_i's are all positive (negative). Our sign convention for double points of the branch locus is chosen for consistency with signs of twisted discs of contact as defined in the paper [12], see Figure 22.
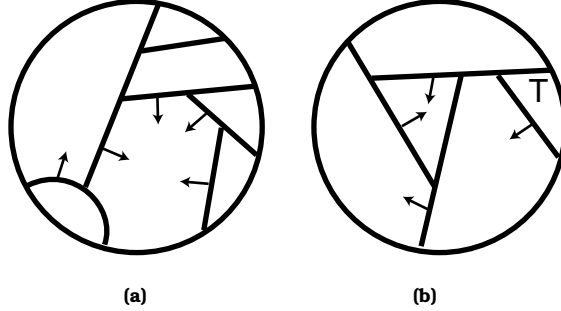

(a)   (b)

Figure 21: T-patterns on discs.

positive double point   negative double point
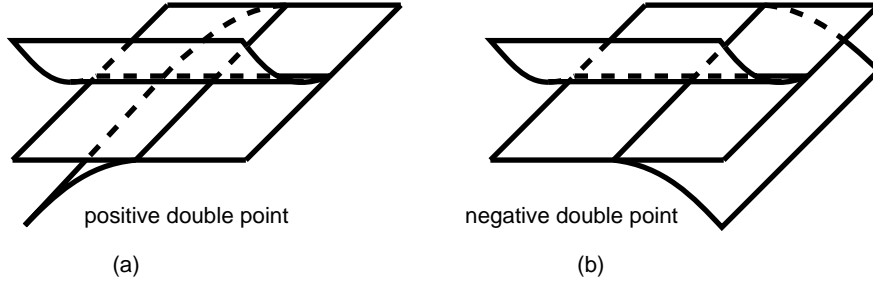
(a)   (b)

Figure 22: The sign of a double point of the branch locus.

We say that B has an *immersed twisted disc of contact* if there exists an immersion $i : E \to B$ of a disc E such that the induced T-pattern on one side has a twisted disc of contact.

Remark 13.2. Deciding wheather a given branched surface has an immersed twisted disc of contact is in general not easy. There are some obvious cases when nonexistence of these objects is obvious. One of them is when a branched surface has no double points in the branched locus.

In our paper [12] we show how to decide whether a given branched surface admits so-called immersed twisted surfaces of contact. The nonexistence of these more general objects clearly implies the nonexistence of immersed twisted discs of contact.

Theorem 13.3. *Suppose a branched surface $B \hookrightarrow M$ carries a positive pure contamination ξ. Suppose also that the branch locus of B consists of essential closed curves, $\partial_h V(B)$ contains no sphere component, and B has no immersed twisted disc of contact. Then ξ is internally universally tight.*

To prove above theorem we need a technical result which allows a closer study of overtwisting discs entirely contained in $int(V(B))$. In this result we use a notion of a branched surface with boundary. We assume that the boundary of a branched surface is a train track transverse to the branch locus (and it avoids double points of the branch locus). The horizontal boundary of a neighbourhood of such a branched surface consists of components that are surfaces with corners, where the corner points separate segments of the cusp locus and boundary segments (i.e. segments that are projected by π to the boundary of the branched surface), see Figure 1.

Proposition 13.4. *Let ξ be a pure contamination carried by a branched surface B, and suppose that $D \subset int(V(B))$ is an overtwisting disc for ξ. Then there is a compact branched surface C with boundary and a map $f : C \to B$ such that:*

(1) C is simply connected;

(2) f is an immersion and when restricted to $int(C)$ it is a local homeomorphism;

(3) if $\bar{\xi}$ is a pullback of ξ to $V(C)$ via f then D pulls back to an overtwisting disc $\bar{D} \subset int(V(C))$ for $\bar{\xi}$.

Proof. Put the overtwisting disc D in general position with respect to the projection map $\pi : V(B) \to B$. By this we mean that the restriction of π to D is generic. More precisely, the projection map has folds and Whitney folds, see Figure 23, and the family of lines in B consisting of the cusp locus of B and the projection of the folds of D and of ∂D to B is generic, meaning that self-intersections are at most double and transverse. Consider then a structure of a cell complex on B such that the 1-skeleton $B^{(1)}$ contains:

(1) projections of folds of D;

(2) the projection of ∂D;

(3) the branch locus of B.

Pulling back, we obtain a cell structure for D.
In a cell structure as above, every cell in D is mapped by π bijectively to its image in B. Modify the cell complex D into a cell complex Q by identifying recursively pairs of 2-cells that intersect (at least at a vertex) and have the same
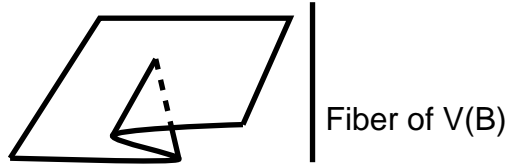
Figure 23: Whitney fold.

projection to B. Use identifications that are compatible with projections to B. Repeat this in arbitrary order, till no pair of 2-cells as above is left in the final quotient. Observe that the resulting complex Q admits the induced from π map $g : Q \to B$ which is an immersion (i.e. a local embedding). Observe also that the identifications we perform produce no new generators in the fundamental groups of succesive quotients, and thus the final complex Q is simply connected, though not necessarily contractible.

To get a branched surface C as required, we want to enlarge Q slightly as follows. If $g : Q \to B$ happens to be an embedding, we take C to be a small regular closed neighbourhood of $g(Q)$ in B. If g is not an embedding, we construct a small regular neighbourhood of the image locally, so that we get kind of a regular neighbourhood of an immersed object. In either of the cases we obtain a branched surface C with boundary for which Q is a deformation retraction. We obtain also an extension $f : C \to B$ of g which, by the local nature of the construction, is an immersion and, at each interior point of C, a local homeomorphisms.

It remains to prove part (3) of the proposition. View the quotient map from D to Q as a map $q : D \to C$, and observe that, viewing D as a subset of $V(B)$, we have $f \circ q = \pi$. On the other hand, we can identify the fibred neighbourhood $V(C)$ with the space

$$\{(x, t) \in C \times V(B) : f(x) = \pi(t)\}.$$

Then the pull-back $\bar{D}$ of D in $V(C)$ is the image of the map $\varphi : D \to V(C)$ defined by $\varphi(u) = (q(u), u)$. $\qquad\square$

One more property of a branched surface C as above is described in the following.

Lemma 13.5. *If C is an orientable simply connected compact branched surface with boundary then the boundary $\partial V(C)$ consists of 2-spheres.*

Proof. It is known that if M is a compact orientable 3-manifold then the image of the boundary map $H_2(M, \partial M) \to H_1(\partial M)$ has rank equal to one half of the rank of $H_1(\partial M)$. Applying this to $M = V(C)$, which is simply connected, we observe that the rank of the boundary map is zero in this case. Therefore the rank of $H_1(\partial V(C))$ is also zero, hence the lemma. $\qquad\square$

The proof of Theorem 13.3 requires further preliminaries, but we give a brief outline here. We will work with a branched surface C as in Proposition

13.4, containing the overtwisting disc $\bar{D}$ for $\bar{\xi}$ in $int(V(C))$. Using a sequence of modifications away from $\bar{D}$ we will convert $\bar{\xi}$ to a pure contamination supported by a disc (and still containing $\bar{D}$ in the interior). This will be a contradiction, due to the following fact.

Fact 13.6. *A pure contamination carried by a disc is internally tight.*

The above fact is well known and can be easily deduced, e.g., from Proposition 3.5.6 in [4].
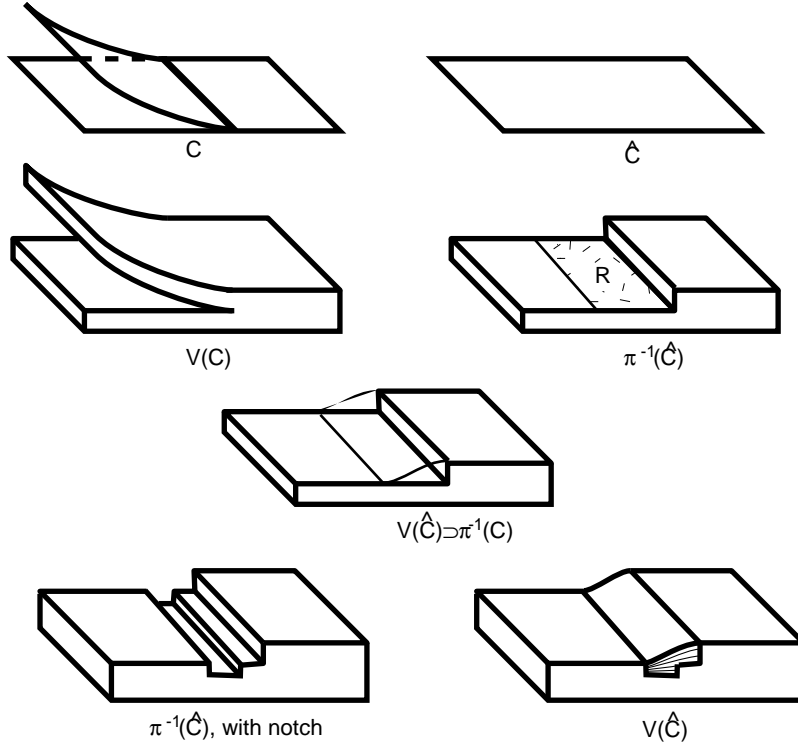

Figure 24: Eliminating a step.

We will now describe a basic modification we will be applying to some parts of $\bar{\xi}$.

Lemma 13.7. *Suppose C is the branched surface shown in Figure 24, and $\hat{C}$ is the disc as shown, obtained from C by deleting one of the sectors. Let α be the cusp arc of $V(C)$ and $R \subset \partial_h V(C)$ a rectangle adjacent to α, as shown in the same figure. Suppose C carries a pure contamination ξ, supported in $V(C)$. Identify $\pi^{-1}(\hat{C})$ with a subset of $V(\hat{C})$ as shown in the figure. Then there is a pure contamination $\hat{\xi}$ carried by $\hat{C}$, supported by $V(\hat{C})$, such that*

(1) $\partial_h V(C) \cap \pi^{-1}(\hat{C}) - R \subset \partial_h V(\hat{C})$;

(2) $\hat{\xi}$ is the same as ξ in $\pi^{-1}(\hat{C})$ except in an arbitrarily small neighborhood in $\pi^{-1}(\hat{C})$ of the rectangle R.

59

Proof. It will be convenient to choose coordinates describing $\pi^{-1}(\hat{C})$. In $\mathbb{R}^3$ $\pi^{-1}(\hat{C})$ will be the set

$$\{(x, y, z) : -2 \leq x \leq 2,\ 0 \leq z \leq 1,\ -4 \leq y \leq 4\}\cup$$
$$\{(x, y, z) : -2 \leq x \leq 2,\ 1 \leq z \leq 2,\ 1 \leq y \leq 4\}.$$

We have a confoliation ξ defined on this set, which is contact except on the horizontal portions of the boundary, where it is integrable. We choose a line segment β at $y = b$, $z = 1$, with $b < 1$ arbitrarily close to the cusp arc α which is at $y = 1$, $z = 1$. The arc β is in the horizontal boundary, where ξ is integrable. We choose a Legendrian γ whose interior is in the interior of $\pi^{-1}(\hat{C})$, and which satisfies $\pi(\gamma) = \pi(\beta)$. The Legendrian γ can be chosen arbitrarily close to the horizontal boundary at $z = 1$, but note that γ need not lie on a level $z = $ constant. Using a σ-deformation, we replace γ by an integrable strip G whose projection in the xy plane contains $b - \epsilon \leq y \leq b + \epsilon$, $-2 \leq x \leq 2$, for some small ϵ, $0 < \epsilon < 1 - b$. We still use ξ to denote the resulting confoliation. Now we remove the portion of the slice $b - \epsilon \leq y \leq b + \epsilon$ above the integral strip G in $\pi^{-1}(\hat{C})$, and we consider the restriction of ξ to this set, which is $\pi^{-1}(\hat{C})$ with a notch cut into it. There is an induced foliation from ξ on all the vertical sides of the notched $\pi^{-1}(\hat{C})$.

Now, we want to insert a foliated set to construct $V(\hat{C})$, with the foliation chosen to give an extension ξ' of ξ as shown in Figure 24, lower right. However, the ξ' obtained is in general not smooth (only continuous) at the vertical sides of the notched $\pi^{-1}(\hat{C})$. Thus, before inserting a foliated part, we need first to modify ξ slightly, so that it becomes a foliation near the vertical sides of the notched $\pi^{-1}(\hat{C})$. This can be done using appropriate charts and modifying slope functions. After this preparation, and after inserting a foliated part as before, we get a smooth ξ' in $V(\hat{C})$ which satisfies the assertions, except that it is not pure.

As a final step, we use Lemma 12.1 to purify ξ' and obtain a pure contamination $\hat{\xi}$ supported in $V(\hat{C})$ and carried by $\hat{C}$. $\qquad\square$

Proof of Theorem 13.3. Suppose $B \hookrightarrow M$ is as in the statement of the theorem, with ξ a pure contamination carried by B and supported in $V(B)$. Then the universal cover $\tilde{B}$ has a neighborhood $\tilde{V}(B)$ which supports a pure contamination $\tilde{\xi}$. We will work in $\tilde{V}(B)$. Assuming there is an overtwisting disc D for $\tilde{\xi}$ in $\mathrm{int}(\tilde{V}(B))$, we will obtain a contradiction.

Let $C, \bar{\xi}$ and $\bar{D}$ be objects as in Proposition 13.4, for the branched surface $\tilde{B}$, the pure contamination $\tilde{\xi}$ and the overtwisting disc D. Our goal is to perform modifications on C and $\bar{\xi}$ to put $\bar{D}$ into a product $E \times I$, as an overtwisting disc of a modified $\bar{\xi}$, which is a pure contamination carried by the disc E and supported by $V(E) = E \times I$.

By the assumptions of this theorem, the branched surface C has no sphere components of $\partial_h V(C)$, no immersed twisted discs of contact, and no closed curve in the cusp locus of $V(C)$. The latter follows from the fact that each cusp curve in the initial branched surface B is essential.

We make the following claim:

Claim: Suppose C is as in the conclusion of Proposition 13.4. Suppose also that C has no sphere components of $\partial_h V(C)$, no immersed twisted discs of contact, and no closed curve in the cusp locus of $V(C)$. Then there is a branch arc α in C cutting a half-disc sector H from C such that the sense of branching is outward from H and $\partial H - \alpha \subset \partial C$.

We will call such an arc α *innermost.* Observe that the cut half-disc H can be lifted in a unique way to a half-disc component of $\partial_h V(C)$ which is adjacent to α viewed as a cusp curve in $V(C)$. Accepting the claim for the moment, we perform a pinching of $V(C)$ identifying a portion of H with a small neighborhood in $\partial_h V(C) - H$ of α (regarded as a cusp curve of $V(C)$). Clearly there is such a pinching which ensures that after the pinching, H is smaller, and D is disjoint from $\pi^{-1}(H)$. We perform a σ-deformation to purify ξ. Next, we let $\hat{C}$ be the branched surface consisting of all sectors of C except H. Then $\pi^{-1}(\hat{C})$ has a step corresponding to the branch curve α. Using Lemma 13.7 we modify ξ to obtain $\hat{\xi}$ carried by $\hat{C}$, and $\hat{C}$ has fewer branch curves than C. Note that all modifications could be done without disturbing the overtwisting disc D. Now we can repeat the process after first renaming as C the branched surface $\hat{C}$, and renaming as ξ the pure contamination $\hat{\xi}$. The new C has the properties of the old one; in particular, it has no immersed twisted disc of contact, since the branch curves of the new C are also branch curves of the old C.

After finitely many modifications, we must obtain a branched surface C without branch arcs, which is still simply connected. Thus C is the required disc E, and we have our overtwisting disc in $\mathrm{int}(V(E))$, contradicting Fact 13.6.

It remains to prove the claim used in the above argument, and we devote the remainder of the section to this aim. $\qquad\square$

Under the hypotheses on C of the Claim we prove the following

Lemma 13.8. *For each component S of $\partial V(C)$ there there are at least two disc components Δ of the horizontal boundary $\partial_h V(C)$, contained in S, with $\partial\Delta$ disjoint from the cusp locus of $V(C)$.*

Proof. It follows from Lemma 13.5 that each component S of $\partial V(C)$ is a 2-sphere. Moreover, no such S is a component of $\partial_h V(C)$. Thus on each S we have a nonempty train-track neighbourhood corresponding to $\pi^{-1}(\partial C) \subset \partial V(C)$ and a family of cusp segments of $V(C)$ (possibly empty) connecting cusp points of the train-track neighbourhood in pairs, passing through the complement of the train-track neighbourhood. (Recall that there are no closed cusp curves in $V(C)$.)

Consider the complementary regions of the train-track neighbourhood in S. Since S is the 2-sphere, an Euler characteristic argument shows that there are at least two complementary regions Δ being discs with smooth boundary (i.e. $\partial\Delta$ contains no cusp point of the train-track neighbourhood). Since there are thus no cusp curves inside Δ, it follows that each of Δ is a component of the horizontal boundary $\partial_h V(C)$ as required. $\qquad\square$

We will call a component Δ as in the above lemma a 0-*gon component* in $\partial_h V(C)$.

We say a region in the complement of a T-pattern (on a component of $\partial_h V(B)$) is *minimal* if all branching at the boundary is outward.

Lemma 13.9. *If a T-pattern on a disc contains no twisted discs of contact, then the T-pattern contains a minimal region.*

Proof. Suppose that the T-pattern on a disc D contains no minimal region. Start in any complementary region R_1 of the T-pattern; there must be at least one transverse orientation on an edge of the T-pattern pointing into R_1, say from R_2. Similarly, there must be at least one transverse orientation on an edge of the T-pattern pointing into R_2, from R_3 say, etc. We conclude that there is a simple closed oriented path γ intersecting edges of the T-pattern with orientation agreeing with the transverse orientation on the edge, and intersecting the T-pattern non-trivially. Now γ bounds a T-pattern on a smaller disc H bounded by γ. Choose any endpoint of the T-pattern on ∂H and suppose that the transverse orientation points to the right. Follow the edge starting at this point until it ends at a T-junction, i.e. ends as the upstroke of the T-junction. Turn right at this T-junction and follow the new edge until it ends at another T-junction, etc. We must obtain a cycle of edges, each starting at a horizontal stroke of a T and extending to the right. Since the T-pattern is on a disc, the cycle represents a twisted disc of contact. We obtain a twisted disc of contact of the opposite sign if the transverse orientation at ∂H points to the left. $\qquad\square$

To prove the existence of an innermost arc for C (provided the branch locus in not empty), consider any component S of the boundary $\partial V(C)$. According to Lemma 13.8, there is a 0-gon component Δ in S. The T-pattern on Δ is then nonempty (otherwise C is already a disc) and contains no twisted disc of contact (otherwise there would be an immersed twisted disc of contact in the initial branched surface B). It follows from Lemma 13.9 that this T-pattern contains a minimal region R. If R is a half-disc (i.e. contains only one arc of the T-pattern in its boundary), then the segment of T-pattern in its boundary is an innermost arc as required. If not (i.e. if R has several arcs of the T-pattern in its boundary), then we use the following inductive argument.

Cut C along an arc β corresponding to an arc in the minimal region R connecting two distinct segments in ∂R not in the T-pattern, see Figure 25. Observe that, since C is simply connected, β disconnects C yielding two simply connected branched surfaces (after necessary smoothing near the ends of β) that are simpler than C (they contain fewer segments in the cusp locus). Choose either of the two branched surfaces and denote it by C'. Observe that after cutting along β the part of the 0-gon component Δ contained in $\partial_h V(C')$ is a 0-gon component Δ_0. Denote by Δ' a different 0-gon component in $\partial_h V(C')$, which always exists due to Lemma 13.8. Note then that $\partial\Delta'$ does not contain a segment corresponding to the cutting arc β, because the component of $\partial_h V(C')$

other than Δ_0 containing such a segment also contains a segment of the cusp locus in its boundary.
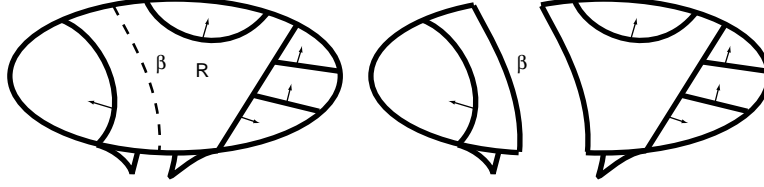

Figure 25: Cutting C on an arc β in a minimal region.

We repeat the previous steps, looking for a minimal region R' of the T-pattern on Δ'. If there is such an R' which is a half-disc, it yields an innermost arc for C', which is also an innermost disc for C (becouse Δ' is far from β). If no such R' is a half-disc, we cut C' as before, along an arc β' in R'. The only difference is that now we choose for C'' the part of C' which does not contain the previous cutting curve β.

After a finite number of steps as above, we must end with an innermost arc for the final branched surface of our sequence. This arc is also an innermost arc for the initial branched surface C, by the choices we make after each cuttting.

This completes the arguments in the proof of Theorem 13.3.

14 Tightness of contaminations

In this section we describe properties of a branched surface B embedded in a 3-manifold M sufficient to guarantee that any pure contamination carried by B is tight. These conditions are based on a similar conditions sufficient for internal tightness given in the previous section. We view the results in this section as promising first steps towards our planned future development of tight pure contaminations as tools in 3-manifold topology.

First we need to recall the definition of incompressibility for a (two-sided) surfaces S properly embedded in a 3-manifold M. A compressing disc for a surface S in a 3-manifold M is a disc D embedded in M, with $D \cap S = \partial D$ and such that ∂D does not bound a disc embedded in S. A surface S properly embedded in M is *incompressible*, if there is no compressing disc for S. We will apply this notion either to closed surfaces S in closed manifolds M, or to the boundaries $S = \partial M$ of compact manifolds M with boundary.

We start with a lemma showing that in some situations it is possible to deduce the tightness of a pure contamination from its internal tightness.

Lemma 14.1. *Let ξ be a pure contamination carried by a branched surface $B \hookrightarrow M$, and suppose that ξ is internally tight. Suppose also that $\partial V(B)$ is an incompressible surface in M and that $\partial_h V(B)$ has no disc components. Then ξ is tight.*

Proof. According to lemma 11.1, we have to exclude existence of two types of overtwisting discs D: the ones with $\partial D \subset \text{int}(V(B))$ and the ones with $\partial D \subset \partial_h V(B)$. By the incompressibility of $\partial V(B)$ in M, if there is an overtwisting disc D for ξ with $\partial D \subset \text{int}(V(B))$, it can be replaced by a disc $D' \subset \text{int}(V(B))$ which coincides with D on a neighbourhood of $\partial D = \partial D'$. Since D' is still an overtwisting disc, this contradicts internal tightness of ξ. Suppose now that D is an overtwisting disc for ξ with $\partial D \subset \partial_h V(B)$. By transversality of D to $\partial_h V(B)$ at ∂D, a neighbourhood of ∂D in D intersects $\partial V(B)$ only at ∂D. Again, by incompressibility of $\partial V(B)$ in M, we can replace D by a disc D' which coincides with D on a neighbourhood of $\partial D = \partial D'$ and which intersects $\partial V(B)$ at $\partial D'$ only. This is a standard argument using the Loop Theorem. Since D is an overtwisting disc, the boundary $\partial D' = \partial D$ does not bound a disc in $\partial_h V(B)$, and because $\partial_h V(B)$ has not disc components, is a compressing disc for $\partial V(B)$ as well, a contradiction. $\square$

Combining the above lemma with Theorem 13.3 we get Corollary 1.16, which can be used to construct examples of tight contaminations.

Example 14.2. The branched surface B shown in Figure 26 can be embedded in $\mathbb{R}^3$, though we have drawn it as an immersed branched surface. By the carrying criterion given in [12], B carries a positive pure contamination ξ.
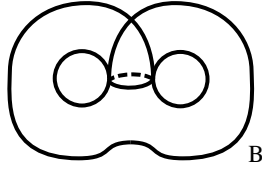

Figure 26: The branched surface B.

Regarding B as embedded in $\mathbb{R}^3$, we have a well-defined neighborhood $V(B)$, and clearly $\partial_h V(B)$ consists of two punctured tori which we label F_+ and F_-. We can glue a 3-manifold with incompressible boundary to $\partial V(B)$ to obtain a closed 3-manifold M.

Observe now that due to Theorem 1.15 the pure contamination ξ is internally tight. To get tightness by applying Lemma 14.1 we need to show that the boundary $\partial_h V(B)$ is incompressible in M.

By construction of M, $\partial V(B)$ is incompressible on the outside of $V(B)$. To prove incompressibility on the inside of $V(B)$, we argue as follows. If α is the branch locus in B, α a closed curve, then $\pi^{-1}(\alpha)$ is an annulus A whose core is the cusp curve $\mathcal{C}(B) \subset V(B)$, cutting A into two annuli A_1 and A_2. Cutting $V(B)$ on A yields a product. If E were a compressing disc for $\partial V(B)$ in $V(B)$, then making E transverse to A and eliminating (trivial) closed curves of intersection, the intersection pattern would consist of arcs. Furthermore, we can assume these arcs are essential in A_1 and A_2, otherwise we could surger to obtain a compressing disc intersecting A in fewer arcs. An innermost arc

would cut a half-disc H from E, with one arc of ∂H in the vertical boundary of the product and one arc in the horizontal boundary of the product. Since this is impossible, we conclude that $\partial V(B)$ is incompressible in M. Thus, by applying Lemma 14.1, we conclude that ξ is tight.

Using Theorem 13.3 it is easy to construct more tight pure contaminations in 3-manifolds, starting with branched surfaces B satisfying the assumptions of this theorem, and then glueing a manifold with incompressible boundary to $\partial V(B)$. The unsatisfactory aspect of these examples, however, is that they can be constructed only in Haken manifolds.

References

[1] B. Aebischer, M. Borer, M. Kälin, Ch. Leuenberger, and H.M. Reimann, *Symplectic geometry*, Progress in Mathematics, vol. 124, Birkhäuser, 1994.

[2] Y. Eliashberg, *Classification of overtwisted contact structures on 3-manifolds*, Invent. Math. **98** (1989), 623–637.

[3] ———, *Legendrian and transversal knots in tight contact manifolds*, Topological Methods in Modern Mathematics, Publish or Perish, 1993, pp. 171–193.

[4] Y. M. Eliashberg and W. P. Thurston, *Confoliations*, University Lecture Series, vol. 13, American Mathematical Society, 1998.

[5] E. Etnyre, *Introductory lectures on contact geometry.*

[6] W. Floyd and U. Oertel, *Incompressible surfaces via branched surfaces*, Topology **23** (1984), 117–125.

[7] D. Gabai and U. Oertel, *Essential laminations in 3-manifolds*, Ann. of Math. **130** (1989), 41–73.

[8] G. Matić K. Honda, W. Kazez, *Convex decomposition theory*, To appear IMRN, Duke U.

[9] ———, *Tight contact structures and taught foliations*, Geometry and Topology **4** (2000), 219–242.

[10] L. Mosher and U. Oertel, *Spaces which are not negatively curved*, Comm. in Anal. and Geom. **6** (1991), 67–140.

[11] U. Oertel, *Measured laminations in 3-manifolds*, Trans. Am. Math. Soc. **305** (1988), no. 2, 531–573.

[12] U. Oertel and J. Świątkowski, *A contamination carrying criterion for branched surfaces*, Preprint, 2001, http://andromeda.rutgers.edu/∼ oertel/.

[13] R. Williams, *Expanding attractors*, Pub. Math. I.H.E.S. **43** (1974), 169–203.